

DISCUSSION PAPER SERIES

IZA DP No. 11661

**The Supply of Skill and Endogenous
Technical Change: Evidence from a
College Expansion Reform**

Pedro Carneiro
Kai Liu
Kjell G. Salvanes

JULY 2018

DISCUSSION PAPER SERIES

IZA DP No. 11661

The Supply of Skill and Endogenous Technical Change: Evidence from a College Expansion Reform

Pedro Carneiro

University College London, IFS, Cemmap and IZA

Kai Liu

University of Cambridge and IZA

Kjell G. Salvanes

Norwegian School of Economics, CELE, FAIR NHH and IZA

JULY 2018

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

The Supply of Skill and Endogenous Technical Change: Evidence from a College Expansion Reform*

We examine the labor market consequences of an exogenous increase in the supply of skilled labor in several cities in Norway, resulting from the construction of new colleges in the 1970s. We find that skilled wages increased as a response, suggesting that along with an increase in the supply there was also an increase in demand for skill. We also show that college openings led to an increase in the productivity of skilled labor and investments in R&D. Our findings are consistent with models of endogenous technical change where an abundance of skilled workers may encourage firms to adopt skill-complementary technologies, leading to an upward-sloping long-run demand for skill.

JEL Classification: J23, J24

Keywords: endogenous technical change, college reform

Corresponding author:

Kai Liu
Faculty of Economics
University of Cambridge
Sidgwick Avenue
Cambridge, CB3 9DD
United Kingdom
E-mail: kai.liu@econ.cam.ac.uk

* We thank Gordon Dahl, Magne Mogstad, Robert Moffitt, Jarle Møen, Jeff Smith, as well as the participants at seminars and conferences for helpful comments and fruitful discussions. This work was supported by the Research Council of Norway through its Centres of Excellence Scheme, FAIR project No 262675. Pedro Carneiro gratefully acknowledges financial support from the Economic and Social Research Council (grant reference RES-589-28-0001 through the Centre for Microdata Methods and Practice), and support of the European Research Council through ERC-2015-CoG-682349. Kjell G. Salvanes and Kai Liu also gratefully acknowledges the support by the Research Council of Norway Grant 202445.

1 Introduction

There are strong links between technological progress and labor markets. Technical change that is skill-biased or complementary to skill (SBTC) is likely to lead to an increase in the skill premium (see, e.g., Katz and Murphy (1992); Autor, Katz, and Krueger (1998); Autor, Katz, and Kearney (2008)), which, in turn, becomes an incentive for individuals to acquire more skill. At the same time, changes in the supply of skill affect the returns to using skill-complementary technologies, and may induce firms to upgrade their technology. The latter mechanism is emphasized in Acemoglu (1998) and Beaudry and Green (2003).

In these papers, an inflow of skilled workers increases returns to using more skill-complementary technologies. If this inflow becomes sufficiently large, firms upgrade their technology. Initially, the skill premium decreases as we move along a downward-sloping demand curve. Once the increase in the supply is large enough for firms to invest in a new technology, the demand for skill shifts outward. As a result, the skill premium and the supply of skilled workers may increase simultaneously.

Our paper provides new evidence that an exogenous shock to the supply of skilled labor induces endogenous technical change. We study data from a college expansion reform in Norway which was rolled out across local labor markets and expanded the supply of college-educated workers, and investigate what happens to wages of skilled workers, the productivity of skilled workers, and R&D investments by firms.

Our paper has three main findings. First, following the opening of a college, both the relative supply of skilled workers and their relative earnings increase simultaneously. In the years immediately after the reform, the increase in the relative supply of skill occurs mainly among young workers (due to the inflow of new university graduates), whereas the increase in the relative earnings of skilled workers occurs mainly among older workers. In the longer run, a college opening induces increases in the relative supply and earnings of skilled workers who were both young and old at the time of the reform.

This finding is consistent with a model where young and old workers are imperfect substitutes (Card and Lemieux, 2001). The earnings of older skilled workers are not very much subject to a downward pressure induced by an increase in the supply of skilled workers, and increase shortly after the opening of a college because of endogenous SBTC. The earnings of skilled young workers are

also affected by endogenous SBTC, but are more subject to downward pressure from the increase in supply. Interestingly, these patterns are much more pronounced following the opening of a STEM college than following the opening of a non-STEM college. Increases in the incentive to invest in new technologies occur mainly in areas where there are increases in the supply of skilled workers in STEM fields.

Second, following the opening of a college, both the supply of skilled workers and their marginal productivity increase simultaneously. The marginal product of skilled and unskilled labor is estimated using plant-level information on output and input factors, ignoring any wage data.

Third, following the opening of a college, firms invest more in R&D (both in terms of expenditure and employment). Together, these three findings suggest that firms responded to the opening of a college, and the resulting increase in the availability of skilled labor, by promoting technical change, either through the adoption of skill-augmenting technologies or changes in organizational form (Acemoglu, 1998; Beaudry and Green, 2003). We argue throughout the paper that, in the period under study, new colleges were not engaging in R&D or innovation activities themselves. They were essentially producing new graduates, so their impact on technical change only occurred indirectly, through the endogenous response of firms to an increase in skill supplies.

Our empirical analysis combines several population-wide and long panel data sets, both containing rich firm-level information on production structure, and individual-level information on demographics and labor market outcomes. Firm-level data span from 1967 to 1990, which includes the population of plants in the manufacturing sector in Norway. Individual-level data combine several administrative registers covering all adult individuals in Norway from the same period. We use the individual data to construct time-series of wage and labor supply by skill groups at the municipality level, which is our baseline definition of a local labor market. We also have information on R&D activities for a subsample of firms, between 1970 and 1985, but not for every year in that interval.

The labor market impacts of college openings are established using a synthetic control method (Abadie, Diamond, and Hainmueller, 2010). There are many fewer municipalities with than without a college opening over the period we study, and this method enables us to find appropriate control municipalities for each municipality with a college opening. Our main results are, however, robust to using a standard difference-in-difference estimator instead. We model the demand of skilled

workers using the data generated from the synthetic control estimator, allowing workers in different age groups to be imperfect substitutes (as in Card and Lemieux (2001)).

Using this model, we quantify the extent to which a college opening induces SBTC. We note that the opening of a college simultaneously affects the supply of skill (through the production of new graduates), and the demand for skilled workers (indirectly, through endogenous SBTC). To separate these two forces, we need an additional assumption. We assume that the impacts of SBTC on the labor market do not take place immediately after the reform, so that only pure supply effects are observed in that period. To be precise, we assume that these impacts do not occur until (at least) two years after the reform. This assumption can be justified if, for example, firms do not invest immediately in response to a college opening, but wait until some of the increase in skill supply materializes. It could also be justified if there are delays in the implementation of a new technology.

The estimation of the impact of the reform on firm-level productivity and firms' R&D activities is performed using versions of more standard difference-in-difference estimators. Our estimation of a firm-level production function is not as amenable to a synthetic control estimator as a reduced-form labor market impact of the reform. In turn, our R&D data do not allow us to use a sufficient number of pre-reform years.

We contribute primarily to the literature examining the rising trend in the college premium (e.g., Katz and Murphy (1992); Berman, Bound, and Machin (1998); Machin and Van Reenen (1998); Card and Lemieux (2001); Autor, Katz, and Kearney (2008)), and the literature on whether the simultaneous increase in the supply of skilled workers and their wages could be due to endogenous SBTC (e.g., Acemoglu (1998, 2007); Beaudry and Green (2003); Blundell, Green, and Jin (2018)). We provide new evidence that endogenous SBTC responds to shocks in the supply of skill led to quantitatively large increases in the skill premium. Our work complements estimates of technology skill complementarity (Akerman, Gaarder, and Mogstad, 2015). If technology and skill are complementary in production, the increase in the abundance of skill induces firms to invest in technology.

This paper is also related to the literature on understanding local economy adjustments to local labor supply shocks. That literature often uses immigration flows in a local labor market as a change in factor supplies. Following a positive shock to the supply of skilled labor, the possible

types of adjustments are through changes in factor prices (by decreasing skilled wages), changes in product mix (by producing a more skill-intensive product mix), and changes in technology (by adopting or spending more to develop skill-biased technologies). The first channel is by no means unimportant, but in light of some evidence that low-skill immigration has little effect on wages, the recent literature has increasingly focused on the latter two channels of adjustments. For example, a number of papers find that most of the adjustment happens through within-industry changes, which they interpreted as changes in production technology (Hanson and Slaughter, 2002; Albrecht, van den Berg, and Vroman, 2009; Lewis, 2011; Dustmann and Glitz, 2015; Peri, 2012).¹ As in these papers, we document that firms adjust their investment in new technologies when faced with a shock to the supply of skilled labor. What is new in our paper is that we document the dynamic impacts of this endogenous technological response on the demand for skilled labor, and, consequentially, on the wages of skilled workers.

Finally, our work is also complementary to recent empirical papers on the reaction of directed technical change to changes in factor supplies (Acemoglu and Finkelstein, 2008; Hanlon, 2015; Bloom, Draca, and Van Reenen, 2016). Our focus is, however, on the labor market.

The paper is organized as follows. Section 2 provides background on the college expansion reform and a description of the data and sample selection procedures. In Section 3, we describe a model of endogenous technology adoption and explain how it guides our empirical work. Section 4 presents reduced-form effects of the reforms on a variety of outcomes at the local labor market level. In Section 5, we provide plant-level evidence on endogenous technical change via estimating production functions. Section 6 presents further evidence that college openings induced firms to invest more in R&D activities. The last section concludes.

¹A related strand of literature studies the impact of high-skilled immigration on innovation and productivity of US firms (see Kerr (2013) for a review). Using state-decadal variation in high-skilled immigration, Hunt and Gauthier-Loiselle (2010) find large increases in innovation following upon immigration. Kerr and Lincoln (2010) find an increasing employment of skilled workers in US firms that experience growth of skilled immigrants with H-1B visas. Peri, Shih, and Sparber (2015) further find city-level productivity increases following from H-1B program expansions in local areas that extensively rely on the program. Moser, Voena, and Waldinger (2014) find that Jewish scientist expellees from Nazi Germany to the US encouraged innovation by attracting new researchers to their fields.

2 Institutional Background and Data

2.1 The College Expansion Reform

The goal of the Parliament when establishing regional colleges in Norway from 1969 onward was to alleviate the increasing problem of capacity at the existing universities.² There was an increasing demand for college education due to a combination of factors, potentially including population growth, changes in the industry composition, and the increased mandatory education, implemented from the late 1950s (e.g., Aakvik, Salvanes, and Vaage (2010), Black, Devereux, and Salvanes (2005)). Figure 1 shows that the overall educational level increased dramatically in Norway from 1960 to 1990 (Statistics Norway, 1994). In 1960, only 4.2 percent of the population above the age of 16 had a college or university degree and 16.3 percent had a high-school diploma, so the ratio of college-educated persons to high school-educated persons was about 0.26. In 1970, these numbers had increased to 6.7 percent and 23.9 percent, respectively, and the ratio of these two quantities rose to 0.28. Then, during the decade we study, the relative supply of college-educated workers rose from 0.28 in 1970 to 0.35 in 1980. This increase in the skill ratio was much larger than what was observed either in the preceding decade (from 1960 to 1970, when this ratio increased from 0.26 to 0.28), or in the following decade (from 1980 to 1990, when this ratio increased from 0.35 to 0.38).

In the mid-1960s, there was strong agreement in the Norwegian Parliament that there was a “national need” to expand the supply of higher education, but there was an understanding that the country did not have sufficient resources to build new universities. A committee appointed by the Government (the Ottosen Committee) was established in 1965, and, in 1966, it proposed to expand the higher education sector by opening regional colleges aiming to provide shorter (two and three years) college education programs (Ottosen-committee, 1966-1970). In its report, the committee proposed to divide the country into 12 educational regions based on the 19 administrative county regions (with some counties merged to a single educational region). Four of the regions already had universities and the committee proposed that the remaining eight regions should each have one regional college.³

²There were only four national universities at the time, located in Oslo, Bergen, Trondheim and Ås.

³In Norway, reports from expert committees are followed up by White Papers from the Government with explicit suggestions to the Parliament to vote on. In the case of the higher education sector over the last several decades, there has been one of these reports from expert committees about every decade (Nyborg, 2007). This is the main procedure determining the total amount of resources allocated to the higher education sector and resource allocations

Following proposals from the Ottosen Committee, in 1968, the Parliament voted for the opening of four new regional colleges.⁴ For the first batch of new colleges, the Parliament initially voted for an experimental period of five years, followed by an evaluation. However, in 1970, the Parliament decided to expand the reform to two more regions,⁵ and, through the 1970s, the establishment of regional colleges was expanded to all the educational regions. There was no discussion in the Parliament’s decision regarding the timing of the reform across these educational regions. Nevertheless, our main identification strategy does not rely on the spatial and temporal variations of the new colleges being exogenous across different educational regions. Instead, as we explain below, we search for reform and non-reform regions with common trends in the main outcomes of interest.

The report from the Ottosen Committee also suggested in which municipality (within each educational region) a new college should be established. Its suggestions were followed up by the Government and voted on by the Parliament. There were three criteria on the geographic location of a college within each educational region. First, the new colleges should be geographically dispersed across the country. Figure 2 shows the geographic location of the new colleges across the country.⁶ This criterion is clearly met for all the new colleges. Second, new colleges had to be established in regions and in municipalities where facilities would be available within a year after the establishment decision took place. Third, the colleges should be used to stimulate growth in regions with stagnation problems.

The two- and three-year programs covered most areas that were available at one of the larger existing universities. In addition, some new programs were developed with two- or three-year durations. For instance, all of the regional colleges had a new program in business administration, which was a new development given that, at the time, business education was only available at one business school in Norway. Over half of the colleges also offered programs in natural sciences, vocational, and technical subjects. It is worth emphasizing that these new colleges aimed to provide college education programs of shorter duration than those offered in the traditional colleges. In

to each college.

⁴Three of them were opened starting from the fall of 1969, located in the counties of Rogaland (in the municipality of Stavanger), Agder (in the municipality of Kristiansand), Møre og Romsdal (dual locations in the municipalities of Volda and Molde). The fourth college was opened in the fall of 1970 in the county Telemark (located in the rural center of Bø).

⁵These are located in the educational region of Hedmark/Oppland (in the municipality of Lillehammer), and one in the region of Nordland (in the municipality of Bodø)

⁶Note that, in this figure, and in our subsequent analysis, we also include three colleges that were built in the same period but were not part of the recommendation of the Ottosen Committee. See the Data Section for details.

the period under study, there was virtually no research output produced by these regional colleges (Ottosen-committee (1966-1970) and Johnsen (1999)). In fact, prior to the 1990s, virtually none of the academic staff in these colleges had a Ph.D. degree (Johnsen, 1999). These colleges were not producing research or innovation, they were producing new graduates.

The size of each new college was also decided by the Parliament based on White Papers from the Government. The Ministry of Education projected the expected demand for study places in each region. These estimates were based on the size of the cohorts in each region, the number of students graduated in the previous years, and the expected demand for higher education. Figure 3 plots the average size of full-time students enrollment, expressed as a percentage of the pre-reform size of college-educated labor (fixed at the 1970 value). Relative to the small size of the pre-existing college-educated labor force, the size of these new colleges potentially had a large impact on the skill compositions in the local labor market: eight years after its opening, the number of full-time enrolled students corresponds to about one-quarter of all college-educated workers in 1970. This suggests that these colleges had a massive impact in the supply of skill in the affected areas.

2.2 Data and Sample Selection

We use both firm- and worker-level data from several sources covering the period 1967–1990. Below, we describe the data we use and the sample construction.⁷

Worker-level data

Our worker-level data come from two sources. The first one contains the data on workers from administrative registers prepared by Statistics Norway. The data cover all Norwegian residents aged 16–74 years old covering the same years as the plant-level data (1967–1990). The variables captured in this dataset include individual demographic information (such as sex and age) and socioeconomic data (such as completed level of schooling, municipality of residence, and annual earnings). For certain male cohorts, we have data on their IQ scores upon entering military service.⁸ In addition

⁷The earliest year of our plant-level and worker-level data begins in 1967. It is potentially interesting to expand the data to analyze the effects of the reform in more recent years. However, following the appointment of a new committee (the Hernes Committee), a new round of reforms was initiated in the early 1990s. By the mid-1990s, regional colleges were consolidated and upgraded to university colleges, where they were given the right to develop research-based degrees, hire professors, and take part in the training of researchers, and to engage in fundamental as well as applied research (Nyborg, 2007). For this reason, we limit our sample period to 1990. During this period, there were no major reforms to the higher education sector.

⁸We explain the use of these data in Appendix Section C.

to the administrative registers, we also use the Norwegian Census, which was conducted in 1960, 1970, and 1980. The census covers the entire population and has additional information on labor market activities (such as industry of employment). A unique personal identifier allows us to follow workers over time and to link the census data with the registry data.

Our wage measure is based on men’s annual labor earnings from the administrative registers.⁹ Annual labor earnings are the sum of pretax labor income (from wages and self-employment) and work-related cash transfers (such as unemployment benefits and short-term sickness benefits). For the period we study, it is not possible to separate the two. We trim the top 0.1% observations with the highest earnings by year, given that top-coding is only performed at very high earnings levels. The Norwegian earnings data have several advantages over those available in most other countries. First, there is no attrition from the original sample because earnings data come from tax records and tax records are in the public domain in Norway. Second, our earnings data pertain to all individuals, and not only to jobs covered by social security.

Education attainment is reported by the educational establishment directly to Statistics Norway, thereby minimizing any measurement error due to misreporting. For every individual, the data record the year of graduation for each level of completed education. Based on this information, we can measure the highest completed level of education for each individual in each calendar year.¹⁰

The sample of individuals being analyzed includes workers aged 20–62 years and whose annual earnings are above the basic amount (1G) required to participate in the national social insurance program.¹¹ In each year, we classify workers into two skill groups. The high-skilled group includes workers who have completed at least some college education. The unskilled group consists of all remaining workers. Our definition of a local labor market is a municipality, which is the smallest administrative unit in Norway.¹² By combining workers’ information on skill levels and municipality of residence, we compute measures of skill composition and wages in a local labor market over time.

⁹Similar to Card and Lemieux (2001), we focus on male wages because there is a large increase in female labor supply over the period under study.

¹⁰The educational establishment data are available starting from 1970. Information on the year of graduation is also left-censored at year 1970. The completed education levels in years 1967–1969 are imputed using the completed education level in 1970.

¹¹Although the mandatory retirement age is 67 years, about 80% of Norwegian workers are entitled to receive early retirement benefits beginning at age 62 years (Bhuller, Mogstad, and Salvanes, 2015). Annual earnings of 1G are the lowest threshold for earning pension points in the national pension scheme. The base amount adjusts for costs of living in each year.

¹²Our definition of the local labor market is consistent with previous empirical work that relies on geographical segmentation of the Norwegian labor market (e.g. Akerman, Gaarder, and Mogstad (2015)).

For instance, the share of high-skilled workers in year t and municipality c is given by the ratio of the number of high-skilled workers over the total size of the labor force residing in municipality c and year t , and the mean wage among high-skilled workers is defined as the average log annual earnings among skilled men residing in municipality c and year t .

Plant-level data

Our main plant-level data are drawn from the Manufacturing Statistics collected annually by Statistics Norway for the period of 1967–1990. The Manufacturing Statistics covers *all* plants in the mining, quarrying, and manufacturing sectors operating during the calendar year in Norway.¹³ The response rate is extremely high because firms are required by law to submit their survey responses.¹⁴ A consistent and unique ID on each establishment allows us to create a panel of plants over this period. We focus on plants in the manufacturing sector with more than five employees and at least 5000 total hours worked in a year. The restriction on size is driven by the fact that complete questionnaires were only collected from plants having at least five employees.¹⁵ The restriction on total hours ensures that the plants in our sample are active in production in any given year.

The firm-level data contains detailed information on output, inputs, and production costs. Using this information, we compute value-added output per firm and year, defined as the gross value of production minus the costs of materials and services.¹⁶

To measure capital stock, we use the fire insurance value of buildings and equipment owned by the firm, and yearly investment flows. The fire insurance value of capital stock is available only from 1974 on. Furthermore, the nature of the fire insurance value means that there is not sufficient variation over time. Therefore, for each plant, we take the fire insurance value as the value of capital in the first year of the panel and impute the value of capital stocks in subsequent years by adding up the value of net investment (by building and equipment separately) in each year and

¹³An establishment is defined as a functional unit that, at a single physical location, is engaged mainly in activities within a specific activity group. A firm doing business in different municipalities is shown as two or more separate establishments in the sample. In the Norwegian context, most plants belong to separate firms (Klette and Griliches, 1996). See Halvorsen, Jensen, and Foyen (1991) for a detailed description.

¹⁴The questionnaires were usually sent out in April/May after the end of the reference year, with a response deadline of four to five weeks. Firms failing to respond to the initial inquiry were sent written follow-up letters for up to six months from the first deadline. Firms that did not respond by then were fined.

¹⁵For small plants with less than five employees, information was extracted from separate administrative registers, which contained fewer variables than the original questionnaire.

¹⁶The value-added output is measured at factor prices, defined as the gross value of production (value of gross output, including subsidies), less the cost of goods and services consumed (excluding VAT) and indirect taxes (except VAT and investment levy).

assuming that the current equipment and buildings depreciate at a constant rate.¹⁷ For plants for which we do not observe the fire insurance value in the first year of the panel (plants which first appeared before 1974 when the fire insurance value was not available), we take the mean fire insurance value by municipality-industry cells in 1974 (separately by buildings and equipment), and use the corresponding cell-specific means as the initial capital stock in the first sampling year of the firm.¹⁸

To measure labor input, we use the total hours of employment for each plant. Unfortunately, the firm-level data do not distinguish labor inputs by skill groups. In addition, for the time period under study, we are not able to link the worker-level data with the plant-level data (this only becomes possible after the 1980s). For this reason, our analysis of worker productivity is conducted at the industry and municipality level because we can observe the skill composition of the labor force at these more aggregated levels by combining worker-level and plant-level data. This is explained in detail in Section 5.

Firm-level R&D data

We also have information on R&D activities for a subsample of firms. During this period, information about R&D is collected from R&D surveys conducted by the Royal Norwegian Council for Scientific and Industrial Research. R&D activities are collected at the line of business level within firms, where a business level includes all plants of the firm for whom the main activity is in the same sector (where sector is defined at a fine level, as explained below). The R&D sample includes mainly manufacturing firms above a certain size-class.¹⁹ We have access to the data starting from 1970, and then biannually from 1975 to 1985. The R&D data can be linked to the plant-level data using a combination of firm and detailed industry identifiers. In Section 6, we discuss our use of the R&D data in detail.

College Reform data

The main source of college reform data is from Ottosen-committee (1966-1970), annual National Budgets (with details on financial support for each college, including the number of students), and

¹⁷The discount rates being used are 0.05 for equipment capital and 0.02 for buildings.

¹⁸The initial value of capital should have little influence on our estimates, given that we conduct our analysis at municipality-industry cell level and control for any permanent differences between cells. See Section 5 for details.

¹⁹The size limits varied among different sectors. The size limits were lower in sectors known to be R&D intensive (down to 10 employees) and higher in sectors with low R&D activity (up to 100 employees). For instance, in the machinery and equipment industries utilized, the R&D surveys have close to full coverage for firms with more than 20 employees. For detailed description of the data, see Statistics Norway (2004) and Møen (2005).

Johnsen (1999), which contains detailed information on the timing, location, programs, and student enrollment of all new regional colleges. Twelve new colleges were built out of the reform initiative in the period we study. We also carefully checked against opening dates of all colleges in Norway and included three additional colleges that were built in the same period but were not part of the recommendation from the Ottosen Committee. The earliest reform year is 1967 and the latest reform year is 1981. The median reform year is 1971.

3 Theoretical Framework

In this section, we review the model of endogenous technology adoption in Acemoglu (2007), and explain how it guides our empirical work. This framework helps us to understand how technology, worker productivity, and wages respond following an increase in the supply of skilled workers. Beaudry and Green (2003) suggest an alternative model of endogenous technology adoption which we also could have adapted for our setting. However, because their model is not focused explicitly on explaining why technological adoption can respond to changes in skill supplies, we chose to discuss instead the model in Acemoglu (2007) which is focused on understanding that issue.

We consider an economy with a set of distinct markets, indexed by i . Consider two types of inputs in production of the final good: S_i is the total amount of skilled labor in market i and U_i is total unskilled-labor supply in market i . For simplicity, factor supplies in each market are assumed to be inelastic, in the sense that they do not respond to changes in factor prices (in this case, wages).

Each market has access to the same set of factor-augmenting technologies θ . For ease of exposition, the set of technologies one can choose from is discrete with two points of support, $\{\theta^a, \theta^b\}$. Suppose that the technology θ^b is more skill-augmenting than technology θ^a .

Each firm in market i chooses factor inputs and the type of technology it wants to adopt. Assuming that the price of the final good is equal to one and that the markets for factor inputs are competitive, the equilibrium in the market can be characterized by one representative firm using aggregate S_i and U_i inputs (Acemoglu, 2007). Equilibrium technology adoption in the market i is given by $\theta^*(S_i, U_i)$, which solves the following problem of the representative firm taking the factor

supplies in the market as given:

$$\max_{\theta} F(S_i, U_i, \theta) = G(S_i, U_i, \theta) - c(\theta)$$

where G is the production function and c is the cost of technology adoption. For simplicity, assume that c is independent of S_i and U_i .

For a market with initial levels of inputs (S_0, U_0) , assume that the initial optimal choice of technology is the least skill-biased one: θ^a . This means that the following condition must hold:

$$c(\theta^b) - c(\theta^a) > G(S_0, U_0, \theta^b) - G(S_0, U_0, \theta^a)$$

This assumption implies that the relative cost of adopting the skilled-biased technology must be large enough to prevent firms from using it.

Now suppose that, at time t , S increases from S_0 to S_1 while the unskilled-labor input is kept fixed at U_0 . As S increases, adopting technology θ^b becomes increasingly attractive because the marginal product of S is higher under θ^b than under θ^a : $\frac{\partial G}{\partial S} |_{\theta^a} < \frac{\partial G}{\partial S} |_{\theta^b}$.

Let S^* be the quantity of skilled-labor input for which the relative cost equals the relative revenue gain of adopting technology θ^b (over θ^a):

$$c(\theta^b) - c(\theta^a) = G(S^*, U_0, \theta^b) - G(S^*, U_0, \theta^a)$$

At S^* (by assumption, $S_0 < S^*$), firms are indifferent between the two technologies.

Therefore, as the economy moves from S_0 to S_1 , the wages of skilled workers change as follows:

$$\Delta w_s = \begin{cases} \frac{\partial G}{\partial S} |_{\theta^a, S=S_1} - \frac{\partial G}{\partial S} |_{\theta^a, S=S_0}, & \text{if } S_1 < S^* \\ \frac{\partial G}{\partial S} |_{\theta^b, S=S_1} - \frac{\partial G}{\partial S} |_{\theta^a, S=S_0}, & \text{if } S_1 \geq S^* \end{cases}$$

If $S_1 < S^*$, $\Delta w_s < 0$, provided the demand for skill is downward sloping. But when $S_1 > S^*$, Δw_s has an ambiguous sign. To see this, decompose Δw_s into a wage change due to supply shift (movement along the demand for skill curve under θ^a) and wage change that is due to technological

upgrading (shift in the demand curve):

$$\Delta w_s = \underbrace{\frac{\partial G}{\partial S} |_{\theta^a, S=S_1} - \frac{\partial G}{\partial S} |_{\theta^a, S=S_0}}_{\text{supply effect}} + \underbrace{\frac{\partial G}{\partial S} |_{\theta^b, S=S_1} - \frac{\partial G}{\partial S} |_{\theta^a, S=S_1}}_{\text{technology effect}} \quad (1)$$

where the supply effect is negative and the technology effect must be positive (because the marginal product of skilled labor is increasing in technology). The net effect could be positive or negative, depending on which effect dominates.

It is not difficult to extend this model to a more dynamic framework where increases in the supply of skilled workers lead to endogenous skill-biased technical change, which, in turn, leads to further increases in the supply of skilled workers. This sort of dynamics may lead to a positive relationship between the quantity of skilled input and the wage of skilled workers. In other words, as discussed in (Acemoglu, 2007), this may lead to a long-run, upward-sloping demand for skill.

4 Worker-level Evidence: Wages, Skill, and Skill-biased Technical Change

We begin this section by documenting positive impacts of a college opening both on the supply of skill and relatively skilled wages in local labor markets. We interpret our estimates using the model of Card and Lemieux (2001), which suggests that for several years following the establishment of a college there is significant skill-biased technology change in the affected labor markets. In Section 5, we provide more direct evidence of endogenous technical change by quantifying the effects of the reform on labor productivity by estimating production functions on firm-level data, as well as estimating the impact of the reform on firms' R&D investments.

4.1 Construction of the Control Group

There are only 15 municipalities that benefited from a college opening and the reform we consider, out of a total of nearly 400 municipalities. In principle, all untreated municipalities can be potential control municipalities, but the danger of proceeding this way is that only a few of them may be

similar to the relatively small set of reform municipalities we consider.²⁰

Therefore, we select comparison municipalities for the control group using the synthetic control estimator developed recently in Abadie, Diamond, and Hainmueller (2010) (hereafter ADH).²¹ For each municipality with a college opening, we use the ADH method to construct an optimal synthetic control group, whose pretreatment trends match that of the treated municipality more closely than those of an arbitrarily-selected set of municipalities. This method is suitable in our setting where a discrete treatment (i.e., a new college) is applied to one unit (i.e., a municipality) and not to others within a large geographic area. The idea is to select control groups based on a set of pre-intervention characteristics Z_{it} which would predict the outcome of interest after the treatment, where Z_{it} includes pre-reform (time-varying) outcome variables (such as the whole history of outcomes), as well as pre-reform (time-invariant) characteristics of the municipality. This procedure provides a vector of municipality-specific weights that minimize the distance between the treated municipalities Z_{it} and the weighted mean of the synthetic control.

In our setting, Z_{it} includes the outcomes measured in each of the five years prior to the treatment, normalized by the outcome in the year of the treatment.²² Z_{it} also includes a set of municipality-level characteristics averaged over pre-reform years, including demographic composition (share of workers aged 20–35 years among the workforce), and skill composition of the labor force (share of high-skilled workers). As a result, for each outcome, the pre-reform trend (the change in the outcome variable in each of the five years prior to the reform), and the skill and age compositions of the labor force in the synthetic control municipality should track closely to those in the treated municipality. Because Z_{it} contains pretreatment outcome variables, a different synthetic control is used for each outcome. To make sure that the control municipality is geographically similar to the treated municipality, we restrict potential control municipalities (donor pool) to be in the same region as the treatment municipality.²³

²⁰Table 1 shows the characteristics between the municipalities with a new college and the remaining municipalities prior to the reform. The log average wages by skill groups are fairly close between treated and remaining municipalities. Relative to the non-treated municipalities, it appears that treated municipalities comprised of a more educated labor force and also experienced faster growth in the skill shares.

²¹As we discuss below, most of our results are robust to using instead a standard difference-in-difference estimator.

²²In cases where the pre-reform period is less than five years, Z_{it} includes the outcomes in all years prior to the treatment, normalized by the outcome in the year of the treatment.

²³We divide the municipalities into four geographical regions as follows: North (Finnmark, Troms, Nordland), Middle (Nord-, Sør-Trøndelag, Møre og Romsdal), West and South (Sogn og Fjordane, Hordaland, Rogaland, Aust-og Vestagder), East (Telemark, Vestfold, Buskerud, Oppland, Hedmark, Oslo, Akershus, Østfold).

4.2 Effects of the Reform on Skill Compositions and Wages

4.2.1 Main Results

Figure 4 presents the effects of the reform on the skill composition of the workforce, and relative earnings of skilled vs unskilled workers, for treatment and control local labor markets (analogous estimates for absolute levels of earnings of skilled and unskilled workers are presented in Figure A1 in the Appendix). Because workers in different age groups are possibly imperfect substitutes in production (Card and Lemieux (2001)), we split the sample into young (aged below 35 years) and old (aged above 35 years) workers, and analyze the impacts of the reform separately for each group.

Workers in the older group may be relatively shielded from the supply effect in equation (1) because the inflow of newly college-educated workers is driven almost exclusively by the young. In turn, workers in the younger age group are affected by both supply and technological effects of the reform.

In each panel of figure 4, the year of the reform for each municipality is normalized to period zero. For each treatment or control municipality, we can compute the difference in the outcome of interest in a given year, relative to the level of that variable in the year of the reform (the level of the outcome in the reform year is also normalized to zero). Each panel in Figure 4 then shows the weighted average of these differences across all 15 sets of treated municipalities (thick line) and the corresponding synthetic controls (dashed line), with weights given by the number of plants in the treated municipality in the given year. The effect of the reform in each year (after year zero) is the difference between the two lines in each panel. Details of our implementation of the ADH procedure are described in the Appendix Section A.

The top-left panel of this figure shows that the reform has a significant impact on the skill composition of young workers, measured by the share of college workers. Compared with the synthetic controls, labor markets with a new college experience an increase in the supply of skilled workers. The gap between the treated group and the synthetic control increases over time (because an additional flow of new graduates is added to the stock of skilled workers each year), reaching nearly four percentage points 10 years after the opening of the college.²⁴ In contrast, the reform

²⁴When we further decompose the unskilled into workers with at least some high school and workers with less than high school, we find that the share of workers with some high school gradually decreases over time in treated

has little impact on the skill composition among workers aged 35 years or more within the first 10 years following the reform.²⁵

The lower panel of Figure 4 shows estimates of the effects of the reform on the relative earnings of skilled workers, by age groups. The relative earnings of young skilled workers are initially similar in both the treatment and control groups, presumably because the supply and technology effects cancel each other. Eventually, the relative earnings of skilled workers in the treatment group increase above those in the control group, perhaps because the technology effect over takes the supply effect.²⁶

Among older workers, the relative earnings of skilled individuals increase substantially following a college opening (this pattern is also seen in the absolute earnings of skilled workers, as shown in Figure A1 in the Appendix). This is consistent with the idea that the earnings of older skilled workers were shielded from any downward pressure induced by an increase in the supply of skill because young and old workers are imperfect substitutes. Their relative earnings increase because the reform increased the demand for skilled labor among older workers, without affecting the supply.

To assess the extent to which our estimates are statistically important, we follow Abadie, Diamond, and Hainmueller (2010) and estimate a series of placebos by iteratively applying the synthetic control method to every municipality in the pool of potential control municipalities. In each iteration, we reassign a treatment from a treated municipality to a control municipality (for details, see the Appendix Section A). This procedure is repeated for each treated municipality so that, for each of them, we obtain an empirical distribution of the estimated gaps between the “treated” municipality and its synthetic control.

In principle, we can calculate p-values for the treatment effects of each treated municipality from the empirical distribution of the gaps implied by the placebos. However, it is simpler to present a single p-value for the treatment effects averaged across all treated municipalities. We

markets. This is, in some sense, expected given that people who are on the margin of going to college are more likely to be affected by the opening up of new colleges and, hence, move from the middle-skilled category to the high-skilled category.

²⁵The share of skilled workers among older workers begins to increase toward the end of the panel, partly because of aging of the cohorts affected by the reform.

²⁶One potential challenge to the interpretation of the wage effects of the young workers is that changes in the relative number of college graduates might affect the relative composition of the pool of college graduates. For instance, after the reform, selection into college- education may be based on ability to a greater extent than mobility costs. To address this concern, we use the IQ information of several cohorts of males from the military draft data. We do not find any evidence that the reform changes the average cognitive ability (proxied by IQ) among college and non-college workers. See the Appendix Section C for details.

begin by randomly drawing 50 placebos (with replacement) for each treated municipality, which we use to compute treatment effects, and then average them across all treated municipalities. We then calculate p-values based on the distribution of the treatment effects from these aggregated placebos.

Figure 5 shows the results of this procedure. The gray lines represent the year-by-year treatment effects for each placebo. The solid black line denotes the treatment effect estimated using the actual data (from Figure 4), with the observed treatment assignment. The implied p-values for each of the actual gaps in each year, i.e., the proportion of placebo gaps larger than the estimated gap, are presented in Appendix Table A1.

There are two variables for which the estimated treatment effects are large and statistically important almost every year after the reform: skill shares among young workers and the relative earnings of skilled older workers. When we consider outcomes many years after the reform, there are statistically significant treatment effects for all outcomes considered in figures 4 and 5, (figure A2 in the Appendix presents the permutation tests for earning levels and Appendix Table A1 shows the corresponding p-values).

Instead of the ADH procedure, we could have used a standard difference-in-differences research design. In Appendix B, we report findings from this exercise where all the untreated municipalities are included as comparisons. The identifying assumption underlying the regression analysis is that the geographic location of the college expansion is not correlated with different underlying trends in local labor-market outcomes across the markets (common trends). As a first check of whether this is a plausible assumption, we examine whether the outcome variable in the treated and control regions have similar trends over time during the pre-reform period. For certain outcomes, the pre-reform trends appear to be different between the treated regions and the remaining areas used as comparison. Therefore, the synthetic control group may be especially helpful in our case for identifying which municipalities should go in the control group. Nevertheless, the effects of the reform on skill and wages across the two age groups are qualitatively similar to the synthetic control estimates.

4.2.2 STEM vs non-STEM Colleges

Out of the 15 new colleges under study, six of them did not have any majors in STEM fields whereas the remaining did have majors in STEM fields. In this section, we ask whether the impacts we estimate are due mainly to openings of STEM colleges because it is plausible that STEM graduates are the ones whose productivity most responds to technical change. Of course, the decision of whether to offer any STEM majors is endogenous and may depend on the existing (pre-reform) local industrial structure, so our estimates have to be interpreted with caution.

Figures 6 and 7 present the estimated effects of reform for STEM and non-STEM colleges for young and old workers, respectively. We find that the opening of both types of colleges led to an increase in the share of young skilled labor in the local labor market, with a stronger effect for STEM colleges. However, the relative earnings responses reported in Figure 4 are driven exclusively by those regions where a STEM college was established. Labor markets where there was an opening of a non-STEM college did not experience substantial changes relative to control labor markets. The results concerning earnings levels (as opposed to relative earnings), which depict an increase in the earnings of skilled workers in areas where there was an opening of a STEM college, are presented in Appendix Figures A3 and A4.

4.3 Separating Supply and Demand Factors

In this section, we use the model in Card and Lemieux (2001) to decompose the differences in trends in skill age-specific wages between reform and non-reform areas (reported in Figure A1) into supply and technology factors. Assume that aggregate output in period t depends on two CES sub-aggregates of skilled (college) and unskilled (non-college) labor:

$$Y_{t(D)} = (\theta_{st(D)} S_{t(D)}^\rho + \theta_{ut(D)} U_{t(D)}^\rho)^{\frac{1}{\rho}}$$

and

$$S_{t(D)} = \left[\sum_j \alpha_j S_{jt(D)}^\eta \right]^{\frac{1}{\eta}}$$

$$U_{t(D)} = \left[\sum_j \beta_j U_{jt(D)}^\eta \right]^{\frac{1}{\eta}}$$

The gross elasticity of substitution between different age groups j with the same level of skill is $\sigma_A = 1/(1 - \eta)$ where $\eta \in (-\infty, 1)$. Different age cohorts of workers are gross substitutes when $\sigma_A > 1$ (or $\eta > 0$), and gross complements when $\sigma_A < 1$ (or $\eta < 0$). If different age groups within a given level of skill are perfect substitutes, η is equal to 1. $\sigma_E = 1/(1 - \rho)$, where $\rho \in (-\infty, 1)$, is the elasticity of substitution between skilled and unskilled workers and substitutes and complements are defined as above. In the CES framework, the values of the elasticities play an important role because they determine how changes in technology and supply of skill by cohorts affect demand and wages. α_j and β_j are efficiency units of skilled and unskilled labor of age group j . Note that this formulation of the CES production function includes factor-augmenting technologies affecting the productivity of workers through the efficiency units of labor.²⁷

Let $D = 1$ denote the treatment group and $D = 0$ denote the synthetic control group. D affects inputs ($S_{jt(D)}$ and $U_{jt(D)}$) and technology ($\theta_{st(D)}$ and $\theta_{ut(D)}$) in the post-reform periods ($t \geq 1$). In period 0 (pre-reform period), the treated group and the synthetic control have identical labor inputs and productivity parameters.

Assuming competitive labor markets (wage equals marginal product of labor), the ratio of wages for skilled and unskilled workers in each age group j is:

$$\log \frac{w_{jt(D)}^s}{w_{jt(D)}^u} = \log\left(\frac{\theta_{st(D)}}{\theta_{ut(D)}}\right) + \log \frac{\beta_j}{\alpha_j} - \frac{1}{\sigma_E} \log \frac{S_{t(D)}}{U_{t(D)}} - \frac{1}{\sigma_A} \left(\log \frac{S_{jt(D)}}{U_{jt(D)}} - \log \frac{S_{t(D)}}{U_{t(D)}}\right) \quad (2)$$

Note that the aggregate supply of skill ($S_{t(T)}$ and $U_{t(T)}$) is unobserved and depends on the parameters in the sub-aggregate CES production function (η , α_j , β_j).

Our main goal is to estimate σ_E , σ_A , and the sequence of $\log\left(\frac{\theta_{st(1)}}{\theta_{ut(1)}}\right) - \log\left(\frac{\theta_{st(0)}}{\theta_{ut(0)}}\right)$. The latter term tells us how endogenous technical change responds to the reform. Although the model of Section 3 is static, it is straightforward to adapt it to a simple dynamic framework, where the choice of technology could affect not only the level of the demand for skill in one period but also its growth over time. This could happen, for example, if the reform leads to a change in the growth of $\frac{S_{t(D)}}{U_{t(D)}}$,

²⁷One notable omission from this model is capital. As emphasized, for example, in Beaudry and Green (2003) and Beaudry, Doms, and Lewis (2010), a decrease in the price of skill could induce an endogenous increase in the capital stock if capital and skilled labor are complements. We introduce capital in the empirical model of Section 5, although, in that specification of the production function, technical change is not allowed to impact the elasticity of substitution between capital and skilled labor. If technical change also makes capital and skill more complementary, as suggested in Beaudry and Green (2003), then we may be overstating the impact of endogenous skill-biased technical change on the relative earnings of skilled workers.

and, therefore, to a change in future endogenous responses to this variable.

There may also exist other reasons, not related to the supply of skill, that led to an impact of college openings on technical change. For example, if new colleges engaged in R&D activities, they could foster an increase in the amount of innovation being produced at any point in time. Below, we discuss why this and other alternatives can be ruled out.

When estimating equation (2), we face two empirical challenges. First, credibly to identify σ_E and σ_A , we need exogenous variation in skill supplies. As argued in the previous section, combining information on college openings with the construction of synthetic controls, we are able to observe arguably exogenous variation in the supply of skill.

Second, any exogenous variation in skill supplies also has an effect on technical change through the channel we emphasize in the paper. Therefore, college openings affect wages through two channels: the direct impact of skill supplies on wages through σ_E and σ_A , and the indirect impact of skill supplies on wages through $\frac{\theta_{st(D)}}{\theta_{ut(D)}}$. Using college openings alone as an exogenous shock is not enough for separately identifying these two mechanisms.

To make progress on this issue, one needs additional assumptions. A reasonable possibility is to assume that $\frac{\theta_{st(D)}}{\theta_{ut(D)}}$ does not respond to D in the years immediately following the reform. This means that $\frac{\theta_{st(0)}}{\theta_{ut(0)}} = \frac{\theta_{st(1)}}{\theta_{ut(1)}}$ for the first M years after the college opening (although it may obviously vary with t for reasons unrelated to the reform, such as exogenous skill-biased technical change). Under this assumption, we can use the (first M) years immediately after the reform to identify σ_E and σ_A in equation (2) for fixed $\frac{\theta_{st(D)}}{\theta_{ut(D)}}$, by relating differences in relative wages to differences in skill shares between reform and non-reform areas. Given σ_E and σ_A , we can use the remaining post-reform years to identify the impact of college openings on $\frac{\theta_{st(D)}}{\theta_{ut(D)}}$.²⁸ For most of our paper, $M = 2$, so this model allows technological progress to influence the labor market shortly after the opening of a college.

To motivate this assumption, notice that if firms are deciding over a discrete set of technologies, then the model discussed in Section 3 (see Beaudry and Green (2003) and Acemoglu (2007)) implies

²⁸Another intuitive idea would be to use the older workers in the years immediately after the reform to identify $\log\left(\frac{\theta_{st(1)}}{\theta_{ut(1)}}\right) - \log\left(\frac{\theta_{st(0)}}{\theta_{ut(0)}}\right)$ because they did not experience increases in $\frac{S_{jt(D)}}{U_{jt(D)}}$ until much later. Given $\log\left(\frac{\theta_{st(1)}}{\theta_{ut(1)}}\right) - \log\left(\frac{\theta_{st(0)}}{\theta_{ut(0)}}\right)$, one could potentially use the younger workers to identify σ_E and σ_A . Of course, this intuition is not quite correct because, even if $\frac{S_{jt(D)}}{U_{jt(D)}}$ does not increase for older workers, their wages are still potentially affected by increases in this variable among the young. The case where this works exactly is when $\sigma_E = \sigma_A$. Under this assumption, age-specific relative wages only depend on age-specific relative supplies.

that an endogenous change in technology only takes place when the supply of skilled workers increases above a certain threshold. Therefore, it is possible that a high enough stock of skilled workers does not build up until a few years after colleges first open, so the endogenous technology responses between reform and non-reform areas may not differ up to that point. Alternatively, if firms are uncertain about the impacts of college openings on skill supplies and require time to learn about it, or if it takes time to implement a new technology, there may not be an immediate reaction of the demand for skilled workers to the opening of a college. Below, we examine the robustness of our estimates to variation in M , the number of years following the reform during which we assume there is no technological response.

We follow Card and Lemieux (2001) to estimate equation (2) in two steps, using the data generated in Figure 4.²⁹ In the first step, σ_A is estimated from a regression of age-group-specific relative wages on age-group-specific relative supplies, age effects, and time effects.

$$\log \frac{w_{jt(D)}^s}{w_{jt(D)}^u} = b_j + d_{tD} - \frac{1}{\sigma_A} \left(\log \frac{S_{jt(D)}}{U_{jt(D)}} \right) + e_{jtD} \quad (3)$$

where b_j and d_t are indicators for age and year effects, respectively. The year effects absorb both the relative technology efficiency between skilled and unskilled labor (and, therefore, it depends on D), as well as any effects of changing aggregate supply. As mentioned above, $\frac{S_{jt(D)}}{U_{jt(D)}}$ could be correlated with e_{jtD} because their unobservable shocks could be driving both demand and supply changes. To explore exogenous changes in supply (and, therefore, identify the demand curve), we explore differences within age (j) and across treatment groups (D):

$$\log \frac{w_{jt(1)}^s}{w_{jt(1)}^u} - \log \frac{w_{jt(0)}^s}{w_{jt(0)}^u} = (d_{t1} - d_{t0}) - \frac{1}{\sigma_A} \left(\log \frac{S_{jt(1)}}{U_{jt(1)}} - \log \frac{S_{jt(0)}}{U_{jt(0)}} \right) + (e_{jt1} - e_{jt0})$$

Given the estimate of σ_A , the efficiency parameters α_j and β_j are estimated via the following

²⁹We estimate this equation in the generated synthetic control data, rather than in the raw data, because the synthetic control procedure allows us to construct a good control group for the treatment firms, which we would not be able to replicate by fitting the model directly to the raw data.

equations:

$$\begin{aligned}\log w_{jt(D)}^s + \frac{1}{\sigma_A} \log S_{jt(D)} &= \gamma_{tD}^s + \log(\alpha_j) + e_{jtD}^s \\ \log w_{jt(D)}^u + \frac{1}{\sigma_A} \log U_{jt(D)} &= \gamma_{tD}^u + \log(\beta_j) + e_{jtD}^u\end{aligned}$$

These equations are derived by equalizing the marginal product of labor with the wage for each combination of age and skill groups. γ_{tD}^s and γ_{tD}^u is a set of year dummies (which vary with treatment), and $\log(\alpha_j)$ and $\log(\beta_j)$ are estimated from the age effects in the above equations. With the estimated σ_A , $\log(\alpha_j)$ and $\log(\beta_j)$, we construct estimates of the aggregate supplies of skilled and unskilled labor in each year for both the treated and synthetic groups.

In the second step of the estimation, we use data from both the treated group and the synthetic group to identify the effects of the college openings on technology change. Based on equation (2), our regression model is:

$$\log \frac{w_{jt(D)}^s}{w_{jt(D)}^u} + \frac{1}{\sigma_A} \left(\log \frac{S_{jt(D)}}{U_{jt(D)}} - \log \frac{S_{t(D)}}{U_{t(D)}} \right) = b_j + \delta_0 t + \delta_1 (t \times P_{t(D)}) + \delta_2 D + \delta_3 P_{t(D)} - \frac{1}{\sigma_E} \log \frac{S_{t(D)}}{U_{t(D)}} + e_{jt(D)} \quad (4)$$

where b_j are age-group dummies, and P is an indicator function that takes value 1 if $t \geq M$ and $D = 1$ and zero elsewhere. Following much of the literature on this topic, the relative technology efficiency, $\log\left(\frac{\theta_{st(T)}}{\theta_{ut(T)}}\right)$, is assumed to follow a linear trend over time and is often interpreted as skill-biased technical change (Katz and Murphy, 1992).³⁰ δ_0 identifies skill-biased technical change in the synthetic control group, whereas δ_1 identifies the impact of the reform on skill-biased technical change.

Table 2 presents the estimated parameters from the second step of the estimation. In column (1), we fix $M = 2$ as our baseline specification, which means that we allow for incremental growth in the relative technology efficiency in treated municipalities two years after the opening of a new college. The interaction between the linear trend and the treatment dummy is positive and significant, which means that the reform leads to an incremental increase in the relative demand

³⁰The linear trend specification is used for parsimony. In theory, we would be able to identify a more flexible version of the trend from our previous assumption. In particular, with the assumption that $\frac{\theta_{st(0)}}{\theta_{ut(0)}} = \frac{\theta_{st(1)}}{\theta_{ut(1)}}$ up to the first L years after the college opening, we can identify σ_E and σ_A from exogenous changes in the supply of skill, obtained from contrasts between areas with and without college openings (because the trend is assumed to be common across these areas in the first L years after the opening of the college).

for high-skilled workers. We estimate that, holding the relative supply fixed, as a consequence of a college opening in the municipality the relative skilled wage increases by an additional of 1.5% per year. The implied elasticity of substitution between college and non-college labor is 1.07.³¹ In subsequent columns, we show that our results are robust to alternative assumptions on the value of M . In column (2), we set $M = 3$ and in column (3), we increase the value of M to 4. Each of these alternatives gives very similar conclusions to our baseline, that there is incremental increase in the relative technology efficiency in treated municipalities following the reform relative to the non-treated municipalities.

We show in Appendix Figure A5 that the estimates of the baseline model of equations (3) and (4) provide an excellent fit to the skill premiums of the treated group and synthetic control group. In Appendix Table A2, we report the estimated parameters from the first step of the estimation. The year dummies show a pattern of steeply rising relative returns. The estimates imply an elasticity of substitution between young and old workers to be slightly below 2, suggesting the importance of considering imperfect substitutability between young and old workers for a given skill group.³²

5 Firm-level Evidence: Structural Estimates of the Production Function

In Section 4, we used labor market data to show that both relative wages and employment of skilled workers increase following the opening of a college. This suggests that the demand for skilled workers increased simultaneously (and, we argue, endogenously) with the increase in the supply of skilled workers.

In this section, we bolster this conclusion by combining quasi-experimental variation from the college reform with estimation of production functions for each industry–municipality combination (constructed by aggregating plant-level data by municipality and industry). Instead of using wage data, we directly estimate the productivity of skilled and unskilled workers, and the impact of the reform on firm output. By using firm-level data, we do not need to rely on the assumption

³¹By comparison, Katz and Murphy (1992) report an estimate of the same parameter equal to 1.4 (using data for both men and women from the US).

³²By comparison, Card and Lemieux (2001) report a higher estimate of the same parameter in the range of 4 to 6 (using data for men from the US). Note that the age groups used in their paper is more disaggregated (five-year age bins defined from ages 26–60 years) than ours.

of competitive labor markets to learn about the impacts of the reform on the productivity of skilled and unskilled labor. Because centralized wage bargaining was rather strong in this period in Norway, fluctuations in productivity may not fully pass onto wages (at least in the short run). This implies that the effects of the reform on labor productivity may be under-estimated if only wage data are used.

We show that both the supply of skill and the marginal product of skilled workers increased in the medium run after the introduction of a new college. This suggests that, in spite of an increase in the amount of skilled labor used by firms, the marginal product of skill increased after college construction. As discussed in Section 3, it is plausible that these results are driven by the adoption of new technologies in firms in response to an increase in the abundance of skilled labor.³³

5.1 Parameters of Interest and Identification

Suppose that the production function for industry j in municipality c and year t takes the following form:

$$Y_{jct} = F(K_{jct}, S_{jct}, U_{jct}, \theta_{ct}) + \mu_{jct} \quad (5)$$

where Y_{jct} is the total real value-added output at factor prices (before taxes). There are three types of inputs used in production: K_{jct} is the total real value of capital stock, S_{jct} is the total employment of skilled workers, and U_{jct} is the total employment of unskilled workers. Inputs, especially S_{jct} and U_{jct} , may depend on \mathbf{d}_{ct} , a vector of current and lagged reform indicators, defined by $\mathbf{d}_{ct} = \{d_{ct}^r\}_{r=0}^R$, where d_{ct}^r is an indicator function $1(t - R_c = \tau)$, R_c is the year of college reform, and τ is the number of years since the reform. θ_{ct} is a vector of skill-augmenting technology parameters which, as mentioned above, may endogenously depend on (S_{jct}, U_{jct}) , and, therefore, on \mathbf{d}_{ct} . μ_{jct} is a productivity shock.

When discussing identification of the parameters of this model, it is useful to rewrite skill-specific

³³In Appendix Figures A6 and A7, we report the synthetic control estimates of the effects of the reform on output per worker at municipality level. The output per worker is measured by taking the mean log value-added per worker over all firms located in the municipality. We find that the reform leads to large and persistent increase in the value-added output per worker, and the increase is exclusively driven by colleges that produce graduates in STEM fields. Note that an increase in output could be driven by changes in technology efficiency or changes in inputs. The estimated production function would allow us to disentangle and quantify the specific channels that drive this increase in output.

labor inputs as a function of total employment (L_{jct}) and skill shares (π_{jct}), where

$$S_{jct} = L_{jct} \times \pi_{jct}$$

$$U_{jct} = L_{jct} \times (1 - \pi_{jct})$$

Therefore, the production function can be equivalently written as:

$$Y_{jct} = F(K_{jct}, L_{jct}, \pi_{jct}, \theta_{ct}) + \mu_{jct}$$

The advantage of this specification relative to the one above is that we summarize the skill content of labor in a single variable, π_{jct} . This way, we can distinguish, in our exposition, endogeneity problems due to L_{jct} (quantity of labor employed) and π_{jct} (skill composition of employment). When we discuss estimation below, we go back to the original specification in equation (5).

Our main objective is to study how the marginal product of skilled and unskilled labor depends on θ , which, in turn, depends on \mathbf{d}_{ct} . We assume that \mathbf{d}_{ct} is exogenous (conditional on covariates, and location and time fixed effects), and, therefore leads to exogenous variation in our main input of interest, π_{jct} , independent of productivity shocks μ_{jct} (and independent of other unobserved input choices). Again, we need to face the problem that π_{jct} has a direct effect on Y_{jct} , but it also has an indirect effect through θ_{ct} . To separate these two channels, we use the same assumption as above: θ_{ct} does not vary with π_{jct} in the first M years immediately following the opening of a college.

We use a control-function approach to deal with the endogeneity of skill shares. To implement it, we start by specifying a reduced-form equation that should be interpreted as an approximation of the determination of the skill shares. This leads to the following econometric model:

$$Y_{jct} = F(K_{jct}, L_{jct}, \pi_{jct}, \mathbf{d}_{ct}^M, \beta) + \mu_{jct}$$

$$\pi_{jct} = G(K_{jct}, L_{jct}, \mathbf{d}_{ct}, \theta) + v_{jct}$$

where $\mathbf{d}_{ct}^M = \{d_{ct}^\tau\}_{\tau=M}^R$. The excluded instruments for π_{jct} are $\mathbf{d}_{ct} \setminus \mathbf{d}_{ct}^M = \{d_{ct}^\tau\}_{\tau=0}^{M-1}$. Assume that

the control function has the following form (analogous to a simple series expansion):

$$E(\mu_{jct}|v_{jct}, K_{jct}, L_{jct}, \mathbf{d}_{ct}) = E(\mu_{jct}|v_{jct}) = \rho_1 v_{jct} + \rho_2 v_{jct}^2 \quad (6)$$

Assuming that K_{jct} and L_{jct} are exogenous, we can implement the following estimator:

$$Y_{jct} = F(K_{jct}, L_{jct}, \pi_{jct}, \mathbf{d}_{ct}^l, \boldsymbol{\beta}) + \rho_1 \hat{v}_{jct} + \rho_2 \hat{v}_{jct}^2 + \omega_{jct}$$

where \hat{v}_{jct} is the estimated residual from the first-stage equation for π_{jct} , which controls for the endogeneity of skill share in the production function. If K_{jct} and L_{jct} are exogenous, then the residual ω_{jct} has zero conditional mean once \hat{v}_{jct} is controlled for.

K_{jct} and L_{jct} are, however, likely to be endogenous as well. Although our main focus in the paper is on π_{ct} and θ_{ct} , we use an additional control function proposed by Levinsohn and Petrin (2003). Our results are robust to the treatment of endogeneity of these other inputs.

Suppose that the demand for intermediate input (material in our case), m_{jct} , depends on the firm's stock of capital and ω_{jct} . Under certain assumptions, LP shows that ω_{jct} , the unobserved productivity shock, can be written as a function of K_{jct} and m_{jct} :

$$\omega_{jct} = \omega_t(m_{jct}, K_{jct})$$

Assuming that capital is additive in the production function, our production function becomes:

$$Y_{jct} = \tilde{F}(L_{jct}, \pi_{jct}, \mathbf{d}_{ct}^l, \boldsymbol{\beta}) + \rho_1 \hat{v}_{jct} + \rho_2 \hat{v}_{jct}^2 + \phi_t(m_{jct}, K_{jct}) + e_{jct}$$

where $\phi_t(m_{jct}, K_{jct}) = f(K_{jct}) + \omega_t(m_{jct}, K_{jct})$. LP use a third-order polynomial to approximate $\phi_t(m_{jct}, K_{jct})$ nonparametrically:

$$\phi(m_{jct}, K_{jct}) = \sum_{p=0}^3 \sum_{q=0}^{3-p} \delta_{pq} K_{jct}^p m_{jct}^q \quad (7)$$

Thus, the unobserved productivity variable ω_{jct} is written as a function of capital and intermediate inputs, both of which are observed in the data. Conditional on $\phi_t(m_{jct}, K_{jct})$, the remaining error

term e_{jct} is not correlated with total employment and capital.³⁴

5.2 Empirical Implementation and Estimation of the Production Function

Given data limitations discussed in Section 2.2, we do not observe skill-specific labor inputs every year at the industry-municipality level. We only observe this variable from the decennial Census data. However, from the administrative register data, we have annual measures of skill shares at the level of the municipality, π_{ct} . It is useful to rewrite skill-specific labor inputs in the production function as a function of observables, π_{ct} and L_{jct} :

$$\begin{aligned} S_{jct} &= L_{jct} Q_{jct} \pi_{ct} \\ U_{jct} &= L_{jct} (1 - Q_{jct} \pi_{ct}) \end{aligned}$$

and

$$Q_{jct} = \left(\frac{S_{jct}}{L_{jct}} \right) / \left(\frac{S_{ct}}{L_{ct}} \right)$$

where S_{jct} and L_{jct} are total skilled labor and total employment in industry j , municipality c , and year t , respectively, and S_{ct} and L_{ct} are total skilled labor and total employment in municipality c and year t , respectively.

We need to make assumptions on how industry-level skill shares grow with municipality-level skill shares. For the results presented here, we assume that Q_{jct} is fixed at its initial value in 1960, which implies that the share of skilled worker in a given industry within a municipality is a constant proportion of the share of skilled workers in the municipality.³⁵ Specifically, we let

$$Q_{jct} = Q_{jc} = \left(\frac{S_{jc60}}{L_{jc60}} \right) / \left(\frac{S_{c60}}{L_{c60}} \right), \forall t$$

One advantage of this assumption is that, by construction, Q_{jc} is exogenous to unobserved productivity shocks (because the first college opening did not take place until the late 1960s).

In the above discussion on identification, we abstracted from any covariates. However, our

³⁴The second step of the LP estimator involves estimating the productivity parameters on capital. For our purpose, the first step is sufficient to recover labor productivity before and after the reform.

³⁵The procedure is described in detail in Appendix D.

empirical strategy requires that we control for a full set of fixed effects for year and municipality because the college expansion reform is plausibly exogenous only conditional on these covariates. Going back to the original specification of the production function, equation (5) (as a function of (S_{jct}, U_{jct}) , not (L_{jct}, π_{jct})), we start by estimating the following empirical model, which accounts for the endogeneity of skill shares (but not the endogeneity of capital or total employment):

$$\begin{aligned}\log(Y_{jct}) &= \beta_{0,jct(D)} + \alpha \log K_{jct} + \beta_{1,ct(D)} \log S_{jct(D)} + \beta_{2,ct(D)} \log U_{jct(D)} + \mu_{jct} \\ &= \beta_{0,jct(D)} + \alpha \log K_{jct} + \beta_{1,ct(D)} \log L_{jct} Q_{jc} \pi_{ct(D)} + \beta_{2,ct(D)} \log L_{jct} (1 - Q_{jc} \pi_{ct(D)}) + \mu_{jct}\end{aligned}\tag{8}$$

$$\pi_{ct(D)} = \theta_1 D_{ct} + \theta_2 D_{ct} \times t_D + \gamma_c^s + \gamma_{k(c)t}^s + v_{ct}\tag{9}$$

where t_D are years since the college opening ($= 0$ in the years up to the college opening, or if there was no college opening in the municipality) and D_{ct} is an indicator function that takes the value 1 if $t_D > 0$.

In this empirical model, we assume that the production function has a standard Cobb–Douglas parameterization. The reform indicator, D , can affect the technology of production in three different ways: the productivity of high-skilled workers, $\beta_{1,ct(D)}$, the productivity of low-skilled workers, $\beta_{2,ct(D)}$, and factor-neutral productivity, $\beta_{0,jct(D)}$. The productivity parameters depend on industry, municipality, time, and lags of the reform according to the following model:

$$\begin{aligned}\beta_{0,jct(D)} &= \gamma_{jc}^0 + \gamma_{jt}^0 + \gamma_{k(c)t}^0 + \delta_1^0(t_D \times P_{t(D)}) + \delta_2^0 P_{t(D)} \\ \beta_{1,ct(D)} &= \gamma_t^1 + \delta_1^1(t_D \times P_{t(D)}) + \delta_2^1 P_{t(D)} \\ \beta_{2,ct(D)} &= \gamma_t^2 + \delta_1^2(t_D \times P_{t(D)}) + \delta_2^2 P_{t(D)}\end{aligned}$$

γ_{jc}^0 are industry by municipality fixed effects, γ_{jt}^0 are industry-by-year effects, and $\gamma_{k(c)t}^0$ are county-by-year effects. γ_d^0 measures the effect of the reform on factor-neutral productivity. γ_t^1 and γ_t^2 are fixed year effects. γ_d^1 and γ_d^2 measure the effects of the reform on the productivity of skilled and unskilled workers, respectively. $P_{t(D)}$ is an indicator function that takes value 1 if $t_D \geq M$, where $M = 2$. The coefficients measuring the impacts of a college opening on factor neutral productivity, the productivity of skilled labor, and the productivity of unskilled labor are (δ_1^0, δ_2^0) , (δ_1^1, δ_2^1) , and

(δ_1^2, δ_2^2) , respectively. This parameterization allows the output elasticities of skilled and unskilled labor to vary with time, location, and with the reform.³⁶

From the first-stage regression (9), we generate estimated residuals \widehat{v}_{jct} . We then include \widehat{v}_{jct} and \widehat{v}_{jct}^2 as controls when estimating the production function (8). This accounts for the endogeneity of the skill composition of employment.

We also estimate another specification of the production function, adding the terms in equation (7) as controls (in addition to polynomials in \widehat{v}_{jct}). This procedure addresses the endogeneity of capital and total labor inputs. Regressions are weighted by the number of plants in each municipality. Standard errors are clustered at the municipality level.³⁷

5.3 Empirical Results

Table 3 reports estimates of equation (8). In column (1), we assume all inputs are exogenous, and, therefore, do not include the terms in either equation (6) or (7). Column (2) adds the control function in equation (6), and column (3) further controls for the terms in equation (7). t_D corresponds to trends in the table, whereas P corresponds to post. Across all specifications, we find that the opening of a college leads to a substantial increase in the estimated output elasticity of skilled labor. At the same time, there is a decrease in the output elasticity among unskilled workers.

If we take, for example, the specification in column 2, five and ten years after the opening of the college, the output elasticity of skilled workers is 4.3 and 7.8 percentage points higher than it would have been in the absence of the reform, respectively. The corresponding numbers for

³⁶We have also tried to allow the coefficient on capital to vary by the reform and year in the same way as the labor inputs. In that model, the estimated effect of the reform on capital productivity is imprecise, and the estimated effects of the reform on productivity of labor inputs are very similar. In addition, we do not find any significant impact of the reform on capital, suggesting that the increase in skilled wages is not driven by an increase in capital. Notice also that our model does not allow technical change to affect capital skill complementarity, as in Beaudry and Green (2003). If technical change also increases capital-skill complementarity, we may be overestimating the direct impact of technical change on the productivity of skilled labor.

³⁷Notice that we estimate equation (8) in the original dataset, as opposed to data generated by a synthetic control estimator, as in Section 4.3. Although this has the disadvantage of being a less reliable research design for determining the causal impact of the reform than the synthetic control method, we saw above that standard difference-in-difference estimates did not produce substantially different results. Our goal here is not to have a reduced-form estimate of the causal impact of the reform on, say, plant output, but to estimate the trajectories of the marginal products of skilled and unskilled labor. The advantage of this procedure is that it provides a more standard treatment of these objects of interest. We could, however, have used a procedure analogous to that in Section 4.3, but it would have been much more cumbersome, given the additional controls we are using here (we would need to estimate a synthetic control estimator for each covariate) and the use of the correction in Levinsohn and Petrin (2003).

unskilled workers are decreases of 3.3 and 7.8 percentage points, five and ten years after the reform, respectively.

Overall, our baseline estimates are consistent with both absolute skill-biased technology change (the productivity of skilled workers increases) and relative skill-biased technology change (the relative productivity between skilled and unskilled workers increases). We also estimate increases in factor-neutral productivity change due to the reform, but the parameters are imprecisely estimated.

Specification Tests

To separate supply and technology effects, we assume that there is no endogenous technological response immediately after the reform. However, under an alternative set of assumptions (arguably less credible), we can allow the technology response to occur immediately after the reform, and we estimate such a model in Appendix Table A6. These alternative assumptions are that: 1) supply shifts are exogenous and 2) there is no endogenous technical change in the control group because the supply shifts are never large enough to justify it. In this case, we can identify the marginal product of skilled and unskilled labor using data from the control group, and the impact of the reform on their productivity using the treatment–control comparison.

Using this model, we find that productivities of both types of labor are unaffected by the reform in the short run. The estimated coefficients on the short-run reform indicators interacted with skilled and unskilled labor are, both individually and jointly, insignificantly different from zero. Therefore, the evidence provides support to our identifying assumption that there are no endogenous technology responses in the short run.

In addition, in Table 4, we show estimates for three other alternative specifications. Columns (1) and (2) show that our results are robust to the inclusion of alternative covariates in $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$. Column (1) drops year fixed effects in $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$. Column (2) includes industry fixed effects in $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$. Column (3) adds municipality group effects in $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$. The coefficients on labor inputs interacted with reform indicators are robust to the above specifications.

5.4 Quantifying the Technology Effects on Wages

To visualize better the quantitative implications of our estimates, in this section, we simulate the predicted marginal product of skilled labor relative to the marginal product of unskilled labor, in

the post-reform years. We decompose changes in this variable into effects of changes in endogenous technology and effects of changes in supply.

From equation (8), it follows that the marginal products of skilled and unskilled labor are given by:

$$\begin{aligned}\log MP_{jct(D)}^s &= \beta_{0,jct(D)} + \alpha \log K_{jct(D)} + \log \beta_{1,ct(D)} + (\beta_{1,ct(D)} - 1) \log S_{jct(D)} + \beta_{2,ct(D)} \log U_{jct(D)} \\ \log MP_{jct(D)}^u &= \beta_{0,jct(D)} + \alpha \log K_{jct(D)} + \log \beta_{2,ct(D)} + \beta_{1,ct(D)} \log S_{jct(D)} + (\beta_{2,ct(D)} - 1) \log U_{jct(D)}\end{aligned}$$

where $\log MP_{jct}^s$ is the log of the predicted marginal product of skilled labor (taking into account the effect of the reform in each year t) and $\log MP_{jct}^u$ is the log of the predicted marginal product of unskilled labor.

In treated municipalities, the predicted change in the log of the relative marginal productivity of skilled labor, $\Delta_{jct(D)} = \log MP_{jct(D)}^s - \log MP_{jct(D)}^u$, due to the reform in year t (after the year of reform) can be decomposed into technology effects and supply effects by the following decomposition:³⁸

$$\begin{aligned}\Delta_{jct(1)} - \Delta_{jct(0)} &= \underbrace{(\log \beta_{1,ct(1)} - \log \beta_{2,ct(1)}) - (\log \beta_{1,ct(0)} - \log \beta_{2,ct(0)})}_{\text{technology effects}} \\ &\quad - \underbrace{[(\log S_{jct(1)} - \log U_{jct(1)}) - (\log S_{jct(0)} - \log U_{jct(0)})]}_{\text{supply effects}}\end{aligned}$$

Figure 8 presents the results of this decomposition. The figure has three lines, representing the technology effect, the supply effect, and the net effect ($\Delta_{jct(1)} - \Delta_{jct(0)}$ in the expression above). Our estimates suggest that the technology effect dominates the supply effect, resulting in an increase in the relative marginal product of skilled labor.

Notice that this predicted increase in $\Delta_{jct(1)} - \Delta_{jct(0)}$ is much larger than the observed increase in the relative wages of skilled workers observed after the opening of a college. For example, 15 years after the reform, our synthetic control estimates are that the relative wages of skilled workers

³⁸ $\log S_{jct(1)} - \log S_{jct(0)}$ and $\log U_{jct(1)} - \log U_{jct(0)}$ are the treatment effects for skilled- and unskilled-labor inputs. We compute these two terms from the estimates of the first-stage regression. Note that, given the construction of the labor inputs discussed previously, the reform affects labor inputs S_{jct} only through local skill share s_{ct} .

are higher in the treatment areas than in the control areas by 10 percentage points, whereas our simulations from the production function estimates suggest that the gap in the marginal product of skilled and unskilled labor grew by 40 percentage points.

As we said above, because centralized wage bargaining is rather strong in Norway, there may be strong deviations from competitive labor market prices, and fluctuations in productivity may not readily translate into fluctuations in wages. Our implied pass-through rate from the marginal product of skilled workers to their wages is about 25%, which is in line with existing estimates from studies done in similar institutional contexts (Margolis and Salvanes, 2001; Barth, Bratsberg, Hægeland, and Raaum, 2012; Akerman, Gaarder, and Mogstad, 2015).³⁹ We cannot rule out the possibility that the estimation method (synthetic control for wages vs difference-in-difference estimator for productivity) also plays a role in this divergence, but we showed above that our estimates of the impact of the reform on wages were robust to the specific method used.

6 Additional Evidence from R&D Activities

We have shown above that the opening of a college leads to changes in the parameters of the relative demand for skill, and changes in the parameters of the production function in the affected municipalities. We interpret our findings so far as evidence of technical change induced by an increase in the supply of skilled workers. In this section, we present direct evidence that college openings induced firms to invest more in R&D activities, which provides further support for the argument put forth in this paper.

In particular, we investigate whether the college expansion reform led to an increase in R&D activities, using firm-level R&D data. The unit of observation in the R&D data is called “*bransjeenhet*,” which consists of all plants of a firm that have their main activity in the same industrial sector. Following Møen (2005), we link the R&D data to our manufacturing plant-level data using a combination of firm and detailed industry identifiers (3-digit). For a firm with a single plant within a single economic activity, this would identify R&D activities at the plant level; for a firm with multiple plants within a single economic activity, we calculate average plant-level values by

³⁹For instance, Akerman, Gaarder, and Mogstad (2015) find that around 20 percent of the increase in marginal productivity of skilled workers (due to firms upgrading their internet technology) is passed through to skilled wages.

dividing the total R&D activities by the number of plants within the same economic activity.⁴⁰

As discussed in Section 2.2, there are fewer observations on R&D than either on wages or firm value-added. The first R&D survey was conducted after the first college was established, and subsequent surveys were conducted only every few years (as opposed to annually). Using these data, we adopt the following empirical specification:

$$\log(Y_{ct}) = \theta_c + \gamma_{s(c)t} + a_1 D_{ct} + \varepsilon_{ct} \quad (10)$$

where Y_{ct} is the sum of all R&D activities among plants located in municipality c and year t ; D_{ct} is an indicator function having value one if the municipality c has a college in year t ; θ_c are fixed municipality effects; and $\gamma_{s(c)t}$ are fixed county–year effects. We consider two measures of R&D activities: i) total costs of performing R&D activities and ii) man-years in performing R&D activities.⁴¹

Table 5 reports our estimates of the parameters of equation (10). Columns (1)–(3) show the effects of the college expansion reform on log total costs of performing R&D activities. The baseline model (column 2) suggests that the reform increases total costs of R&D expenditures in the area by over 80 percent. The estimate is robust to the inclusion of a municipality-specific linear trend (column 3), and replacing county–year fixed effects with year fixed effects (column 1). Columns (4)–(6) show that the reform also has a large positive impact on the man-years performing R&D activities in firms. Including the municipality-specific linear trend makes the estimates less precise (due to the small sample used in this regression), but overall the estimate points to large positive increase in the man-years performing R&D activities (more than doubled) in areas following the reform.

⁴⁰The R&D data only samples large firms above a certain threshold in selected years after 1970 (Section 2.2). Therefore, a smaller number of observations in our plant-level data (2,579 plant-years) are linked to the R&D data. Of the linked observations, there are 55 percent of the firm-activity units with one or two plants and 52 percent of the firm-activity units located in a single municipality.

⁴¹The total costs of R&D include the internal operating costs and the external procurement costs related to R&D activities. This variable is available in all rounds of the R&D survey (1970, 1975, 1977, 1979, 1981, 1983, and 1985). Variables on the man-years in performing R&D activities are only available in 1970, 1975, 1979, and 1985.

7 Conclusion

The leading hypothesis explaining the simultaneous increase in the supply of skill and of the skill premium observed in many countries over the last five years is skill-biased technology change. A related set of papers asks why SBTC was so pronounced during this period, and suggests that it could be an endogenous response to an increase in the relative supply of skilled labor. Whereas predictions from the endogenous technical change hypothesis are shown to be consistent with aggregate supply and wage changes over time, there is little evidence that an exogenous increase in the supply of skill led to additional investments in skill-biased technologies and a resulting increase in the productivity of skilled labor.

In this paper, we examine the consequences of an exogenous increase in the supply of skilled labor in several local labor markets in Norway, resulting from the construction of new colleges in the 1970s. The reform shifted skill compositions of these affected areas over time: regions with a new college had more rapid growth in the share of skilled workers than a set of comparison areas without any new college. We use the spatial and temporal variation in the availability of new colleges across local labor markets as a natural experiment for changes in the local supply of skill. Our empirical analysis draws on several large and long panel datasets containing rich firm-level information on production structure and individual-level information on demographics and labor market outcomes.

We find that local average skilled earnings both relative to unskilled and in levels increased as a response to the new college, which is suggestive of a skill-biased demand shift. Results from our relative labor demand regressions also indicate unobserved technology change favoring college workers relative to high-school workers. Drawing from a large panel of manufacturing firms, our production function estimates also suggest that there are endogenous skill-biased technology investments in response to a college opening because the productivity of high-skilled workers increased after the reform (even after accounting for changes in the capital stock). We interpret our findings using existing models of directed technical change, which predict that an abundance of skilled workers may encourage firms to use more skill-complementary technologies. As a result, the demand for skill may increase, leading to an upward-sloping demand curve in the long run.

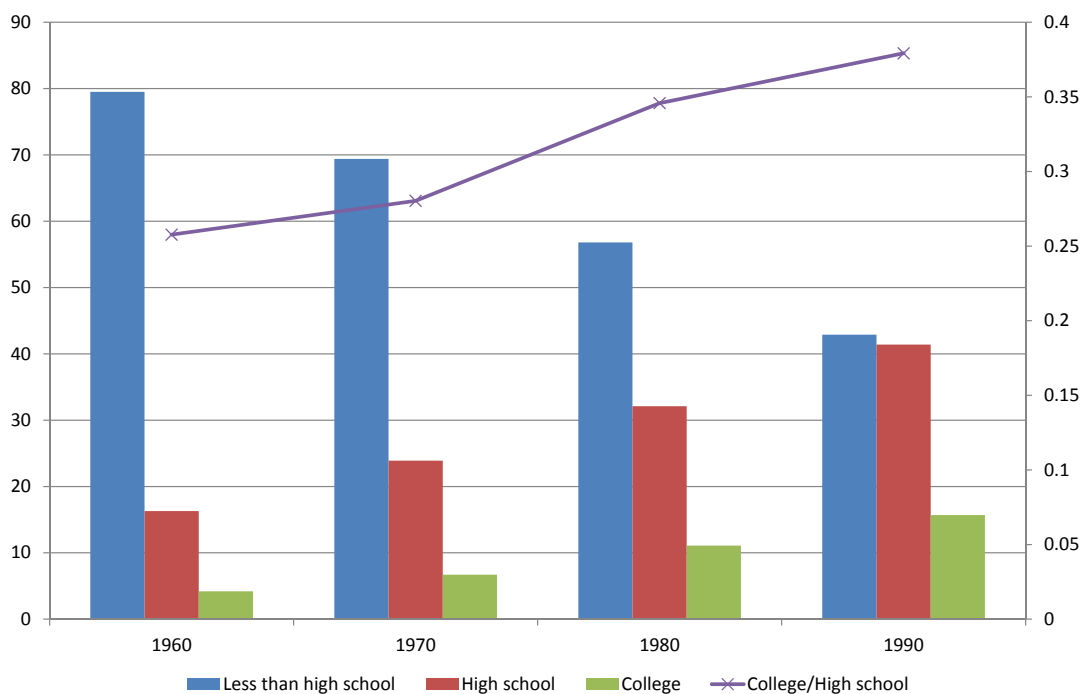
References

- AAKVIK, A., K. G. SALVANES, AND K. VAAGE (2010): “Measuring Heterogeneity in the Returns to Education using an Education Reform,” *European Economic Review*, 54, 483–500.
- ABADIE, A., A. DIAMOND, AND J. HAINMUELLER (2010): “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program,” *Journal of the American Statistical Association*, 105(490), 493–505.
- ACEMOGLU, D. (1998): “Why do new technologies complement skills? Directed technical change and wage inequality,” *Quarterly journal of economics*, pp. 1055–1089.
- (2007): “Equilibrium bias of technology,” *Econometrica*, 75(5), 1371–1409.
- ACEMOGLU, D., AND A. FINKELSTEIN (2008): “Input and technology choices in regulated industries: Evidence from the health care sector,” *Journal of Political Economy*, 116(5), 837–880.
- AKERMAN, A., I. GAARDER, AND M. MOGSTAD (2015): “The skill complementarity of broadband internet,” *The Quarterly Journal of Economics*, 130(4), 1781–1824.
- ALBRECHT, J., G. J. VAN DEN BERG, AND S. VROMAN (2009): “The aggregate labor market effects of the Swedish knowledge lift program,” *Review of Economic Dynamics*, 12(1), 129–146.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2008): “Trends in US wage inequality: Revising the revisionists,” *The Review of Economics and Statistics*, 90(2), 300–323.
- AUTOR, D. H., L. F. KATZ, AND A. B. KRUEGER (1998): “Computing inequality: have computers changed the labor market?,” *The Quarterly Journal of Economics*, 113(4), 1169–1213.
- BARTH, E., B. BRATSBERG, T. HÆGELAND, AND O. RAAUM (2012): “Performance Pay, Union Bargaining and Within-Firm Wage Inequality,” *Oxford bulletin of economics and statistics*, 74(3), 327–362.
- BEAUDRY, P., M. DOMS, AND E. LEWIS (2010): “Should the personal computer be considered a technological revolution? Evidence from US metropolitan areas,” *Journal of Political Economy*, 118(5), 988–1036.
- BEAUDRY, P., AND D. A. GREEN (2003): “Wages and employment in the United States and Germany: What explains the differences?,” *The American Economic Review*, 93(3), 573–602.
- BERMAN, E., J. BOUND, AND S. MACHIN (1998): “Implications of Skill-Biased Technological Change: International Evidence,” *Quarterly Journal of Economics*, pp. 1245–1279.
- BHULLER, M., M. MOGSTAD, AND K. G. SALVANES (2015): “Life cycle earnings, education premiums and internal rates of return,” *Journal of Labor Economics*.

- BLACK, S. E., P. J. DEVEREUX, AND K. G. SALVANES (2005): “Why the Apple Doesn’t Fall Far: Understanding Intergenerational Transmission of Human Capital,” *American Economic Review*, 95, 437–449.
- BLOOM, N., M. DRACA, AND J. VAN REENEN (2016): “Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity,” *The Review of Economic Studies*, 83(1), 87–117.
- BLUNDELL, R., D. GREEN, AND W. JIN (2018): “The UK Education Expansion and Technological Change,” *Institute for Fiscal Studies Working Paper*.
- CARD, D., AND T. LEMIEUX (2001): “Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis,” *The Quarterly Journal of Economics*, 116(2), 705–746.
- CAVALLO, E., S. GALIANI, I. NOY, AND J. PANTANO (2013): “Catastrophic natural disasters and economic growth,” *Review of Economics and Statistics*, 95(5), 1549–1561.
- CRONBACH, L. J. (1964): *Essentials of psychological testing*. Harper and Row.
- DUSTMANN, C., AND A. GLITZ (2015): “How do industries and firms respond to changes in local labor supply?,” *Journal of Labor Economics*, 33(3), 711–750.
- HALVORSEN, R., R. JENSEN, AND F. FOYN (1991): “Dokumentasjon av industristatistikkens tidsseriebase (” Documentation of the Panel Data Base for Manufacturing”),” Discussion paper, Mimeo, Statistics Norway (In Norwegian).
- HANLON, W. W. (2015): “Necessity is the mother of invention: Input supplies and Directed Technical Change,” *Econometrica*, 83(1), 67–100.
- HANSON, G. H., AND M. J. SLAUGHTER (2002): “Labor-market adjustment in open economies: Evidence from US states,” *Journal of International Economics*, 57(1), 3–29.
- HECKMAN, J. J., J. STIXRUD, AND S. URZUA (2006): “The effects of cognitive and non-cognitive abilities on labor market outcomes and social behavior,” *Journal of Labor Economics*, 24(3), 411–482.
- HUNT, J., AND M. GAUTHIER-LOISELLE (2010): “How much does immigration boost innovation?,” *American Economic Journal: Macroeconomics*, 2(2), 31–56.
- JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993): “Earnings losses of displaced workers,” *The American Economic Review*, pp. 685–709.
- JOHNSEN, B. W. (1999): *Fra universitetsvisjon til høyskoleintegrasjon*. Kristiansand: Norwegian Academic Press.

- KATZ, L. F., AND K. M. MURPHY (1992): “Changes in relative wages, 1963–1987: supply and demand factors,” *The quarterly journal of economics*, 107(1), 35–78.
- KERR, W. R. (2013): “US high-skilled immigration, innovation, and entrepreneurship: Empirical approaches and evidence,” Discussion paper, National Bureau of Economic Research.
- KERR, W. R., AND W. F. LINCOLN (2010): “The supply side of innovation: H-1B visa reforms and US ethnic invention,” *Journal of Labor Economics*, 28(3), 473–508.
- KLETTE, T. J., AND Z. GRILICHES (1996): “The inconsistency of common scale estimators when output prices are unobserved and endogenous,” *Journal of Applied Econometrics*, 11(4), 343–361.
- LEVINSOHN, J., AND A. PETRIN (2003): “Estimating production functions using inputs to control for unobservables,” *The Review of Economic Studies*, 70(2), 317–341.
- LEWIS, E. (2011): “Immigration, skill mix, and capital skill complementarity,” *The Quarterly Journal of Economics*, 126(2), 1029–1069.
- MACHIN, S., AND J. VAN REENEN (1998): “Technology and changes in skill structure: evidence from seven OECD countries,” *The Quarterly Journal of Economics*, 113(4), 1215–1244.
- MARGOLIS, D. N., AND K. G. SALVANES (2001): “Do Firms Really Share Rents with Their Workers?,” *IZA Discussion Paper No. 330*.
- MØEN, J. (2005): “Is Mobility of Technical Personnel a Source of R&D Spillovers?,” *Journal of Labor Economics*, 23(1), 81–114.
- MOSER, P., A. VOENA, AND F. WALDINGER (2014): “German Jewish émigrés and US invention,” *American Economic Review*, 104(10), 3222–55.
- NYBORG, P. (2007): “Higher Education in Norway—Fifty years of development,” .
- OTTOSEN-COMMITTEE (1966-1970): *Innstilling om videreutdanning for artianere m.v. fra komiteen til å utrede spørsmål om vidreutdanning for artianere og andre med tilsvarende grunnutdanning*.
- PERI, G. (2012): “The Effect Of Immigration On Productivity: Evidence From US States,” *The Review of Economics and Statistics*, 94(1), 348–358.
- PERI, G., K. SHIH, AND C. SPARBER (2015): “STEM workers, H-1B visas, and productivity in US cities,” *Journal of Labor Economics*, 33(S1), S225–S255.
- STATISTICS NORWAY, N. (1994): “Historisk statistikk 1994,” *Historical statistics*.
- (2004): “Research and development in the Business Enterprise Sector 2001-2002,” *Official Statistics of Norway D*, 301.

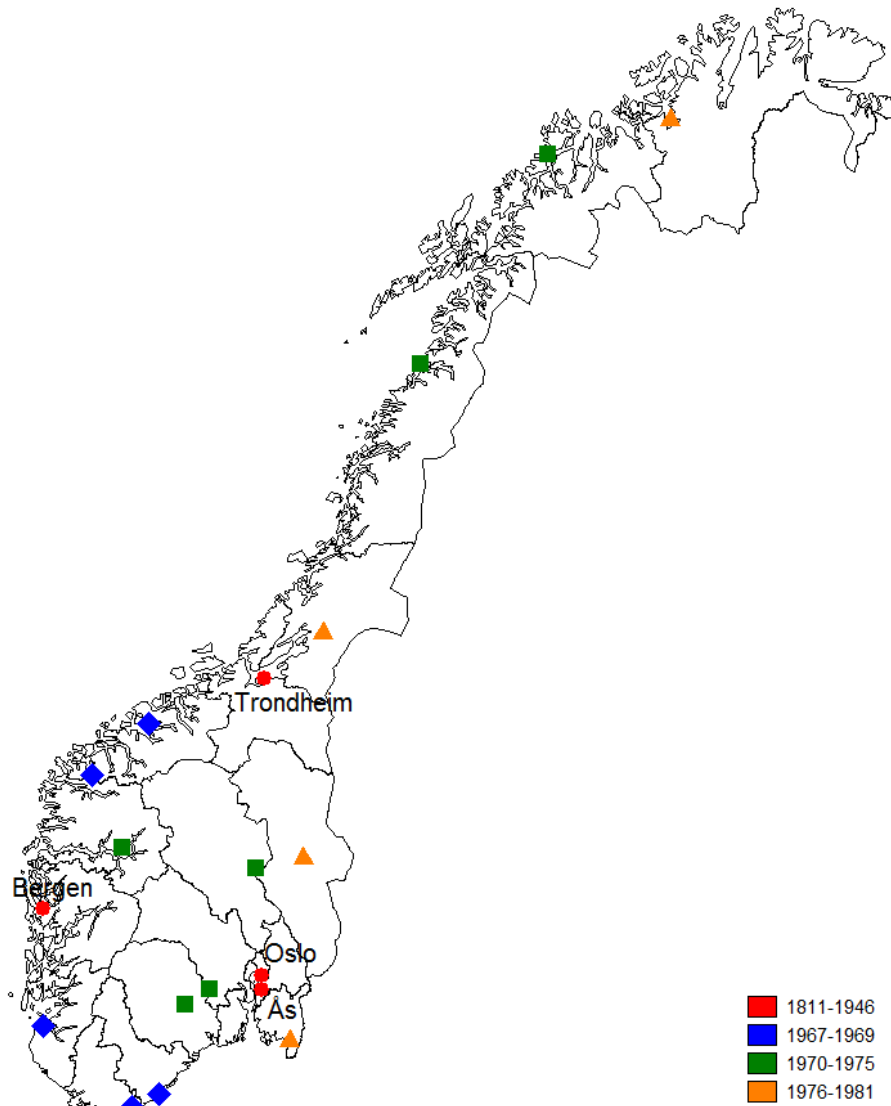
- SUNDET, J. M., D. G. BARLAUG, AND T. M. TORJUSSEN (2004): “The end of the Flynn effect? A study of secular trends in mean intelligence test scores of Norwegian conscripts during half a century,” *Intelligence*, 32(4), 349–362.
- SUNDET, J. M., K. TAMBS, J. R. HARRIS, P. MAGNUS, AND T. M. TORJUSSEN (2005): “Resolving the genetic and environmental sources of the correlation between height and intelligence: A study of nearly 2600 Norwegian male twin pairs,” *Twin Research and Human Genetics*, 8(4), 307–311.
- THRANE, V. C. (1977): “Evnepøving av utskrivingspliktige i Norge 1950-53,” Arbeidsrapport nr. 26, INAS (in Norwegian).



This figure shows the distribution of levels of completed education in the adult population (in percentages; left scale) by year. The line shows the ratio of the college-educated persons over high-school educated persons by year (right scale). Source: Statistics Norway (1994).

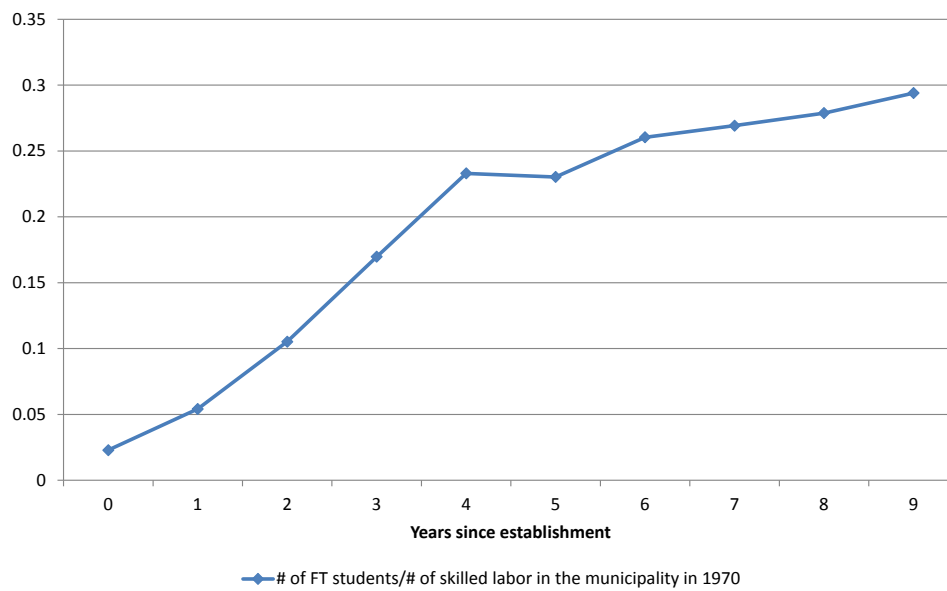
Figure 1: Completed Level of Education for Individuals Aged 16+ (in percentages): 1960–1990

Figure 2: Geographic Locations of New Colleges in Norway: 1967–1985



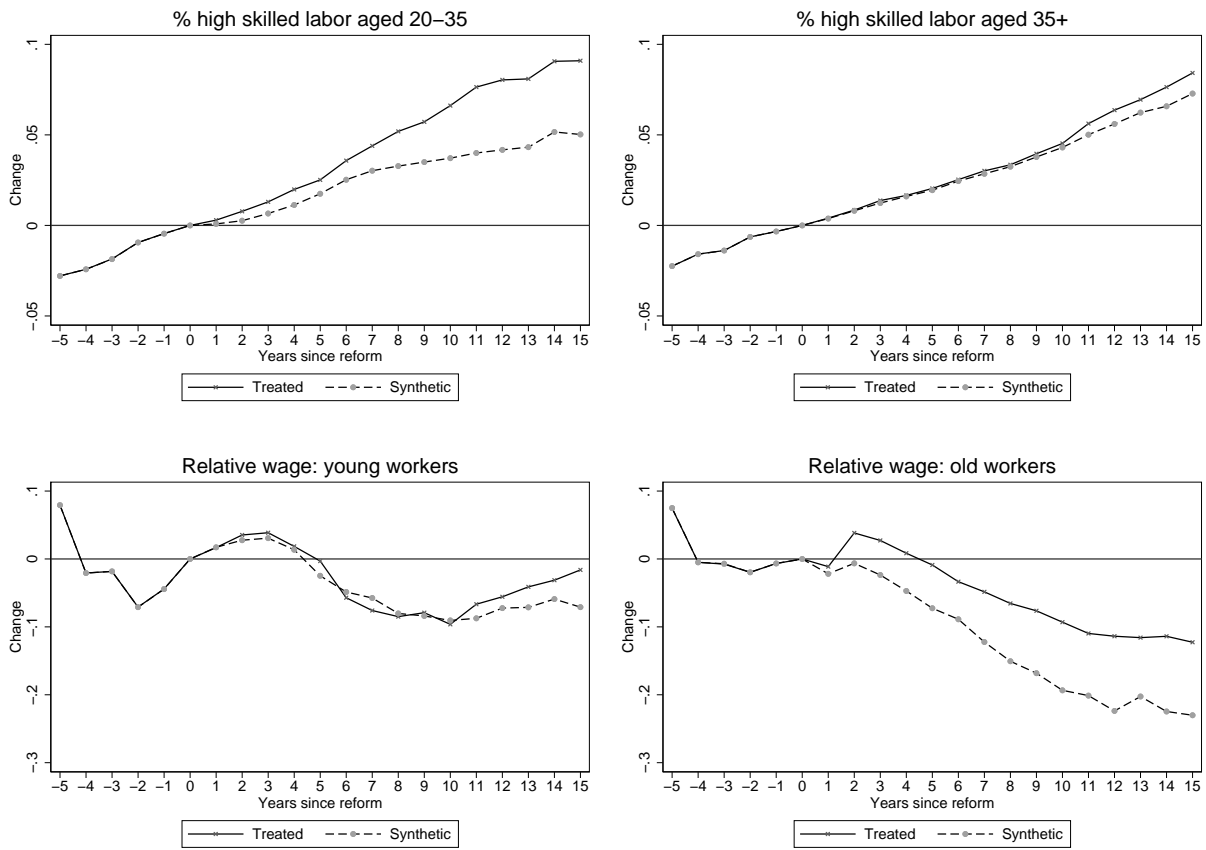
Note: This figure shows the geographic location of colleges across the country. We distinguish the colleges by their establishment year. The four universities established before the reform are labeled in red circles. The remaining colleges shown on this map are labeled based on their year of establishment.

Figure 3: Size of Full-time Students over Size of College-educated Labor, by Years since Establishment



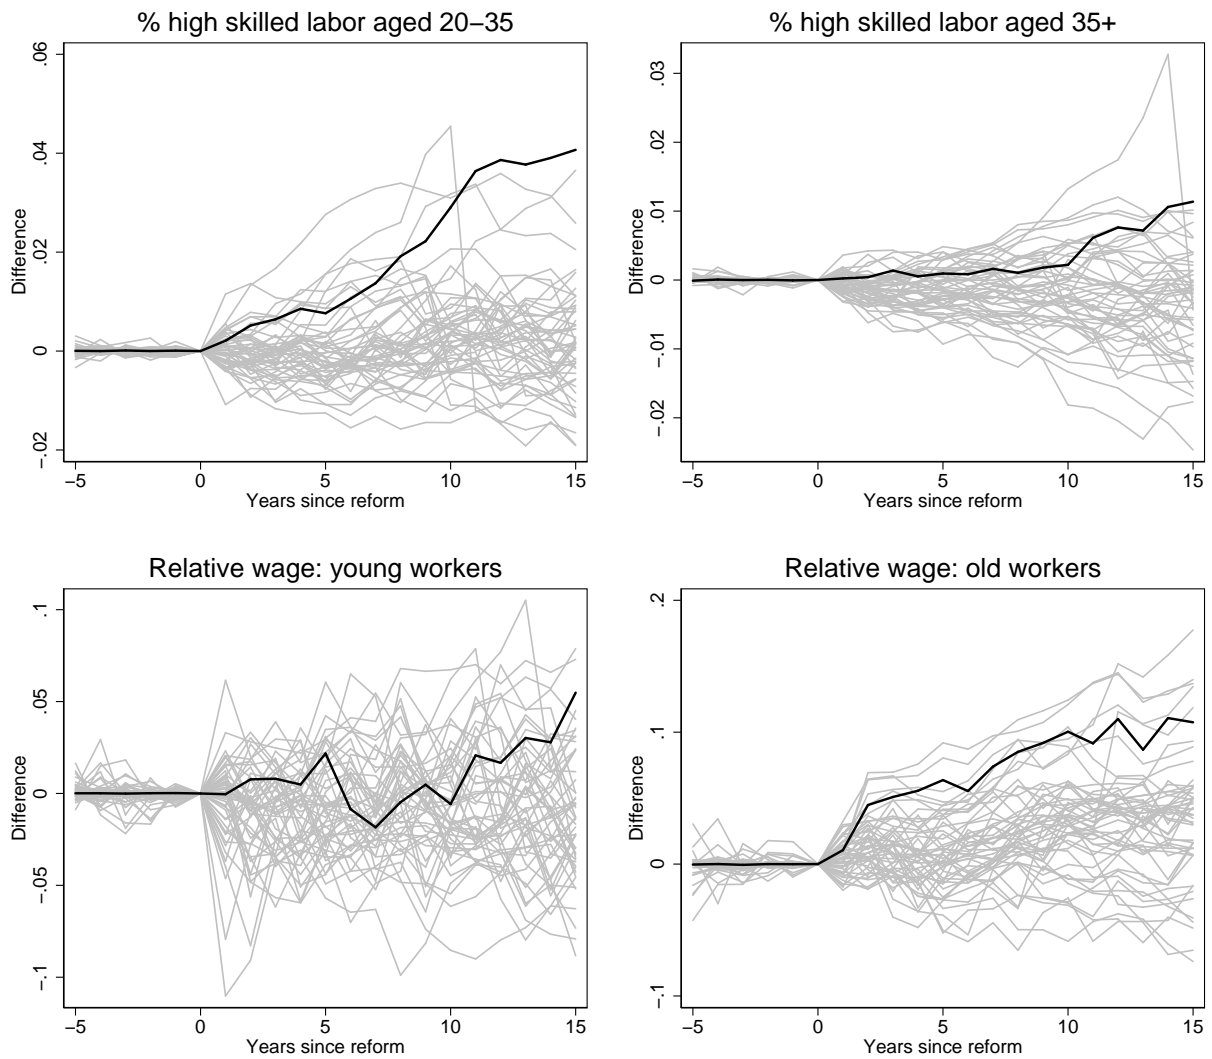
Note: This figure shows the total number of full-time students enrolled in the seven new colleges (built between 1969 and 1971) in each year after establishment as a percentage of the total number of college-educated labor in municipalities of the new colleges in 1970. Year 0 is normalized as the establishment year. Data source: Norwegian Government National Budget, 1969–1980.

Figure 4: The Effects of the Reform on Relative Wages and Skill Compositions



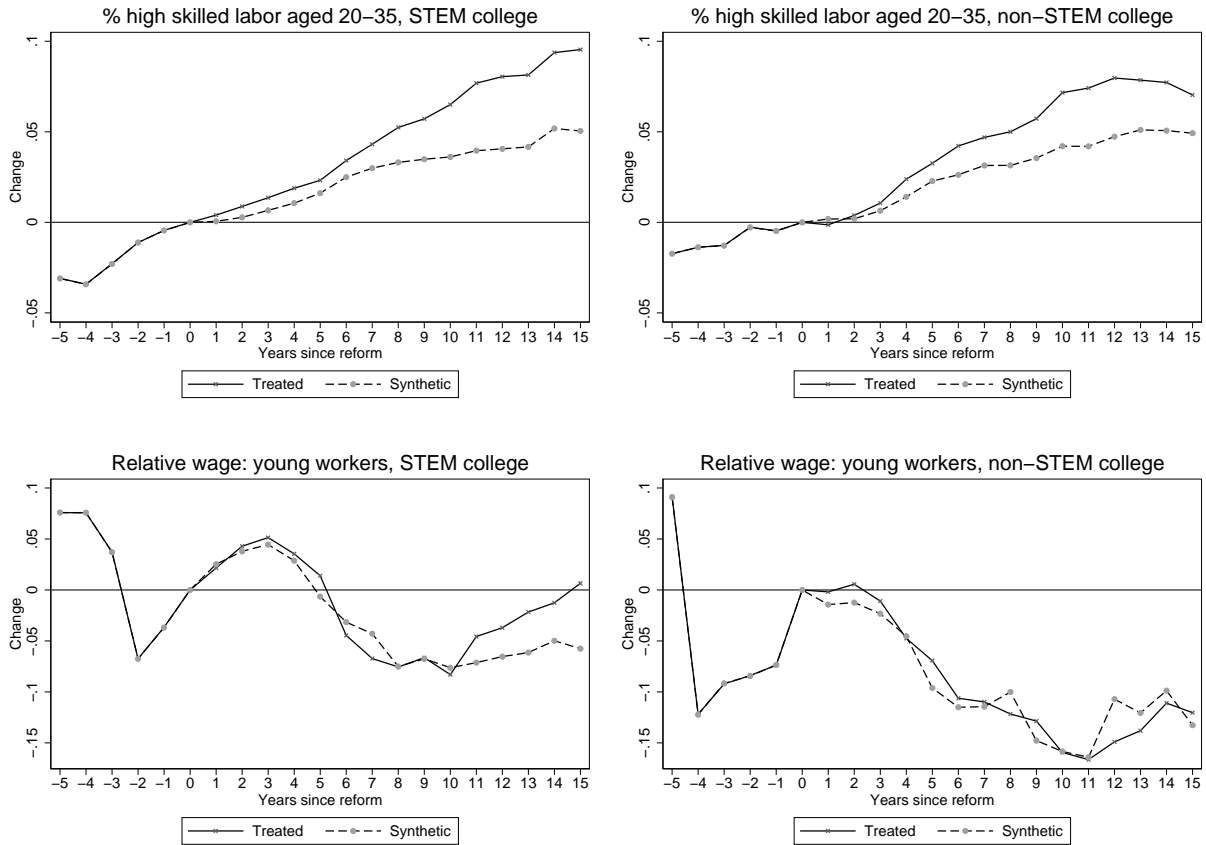
Note: This figure presents the synthetic control estimates on skill composition and relative wages of the workforce. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). Each graph reports the weighted average of all treated municipalities and the corresponding synthetic controls, with weight given by the number of plants in the treated municipality.

Figure 5: The Effects of the Reform on Relative Wages and Skill Compositions: Placebo Tests



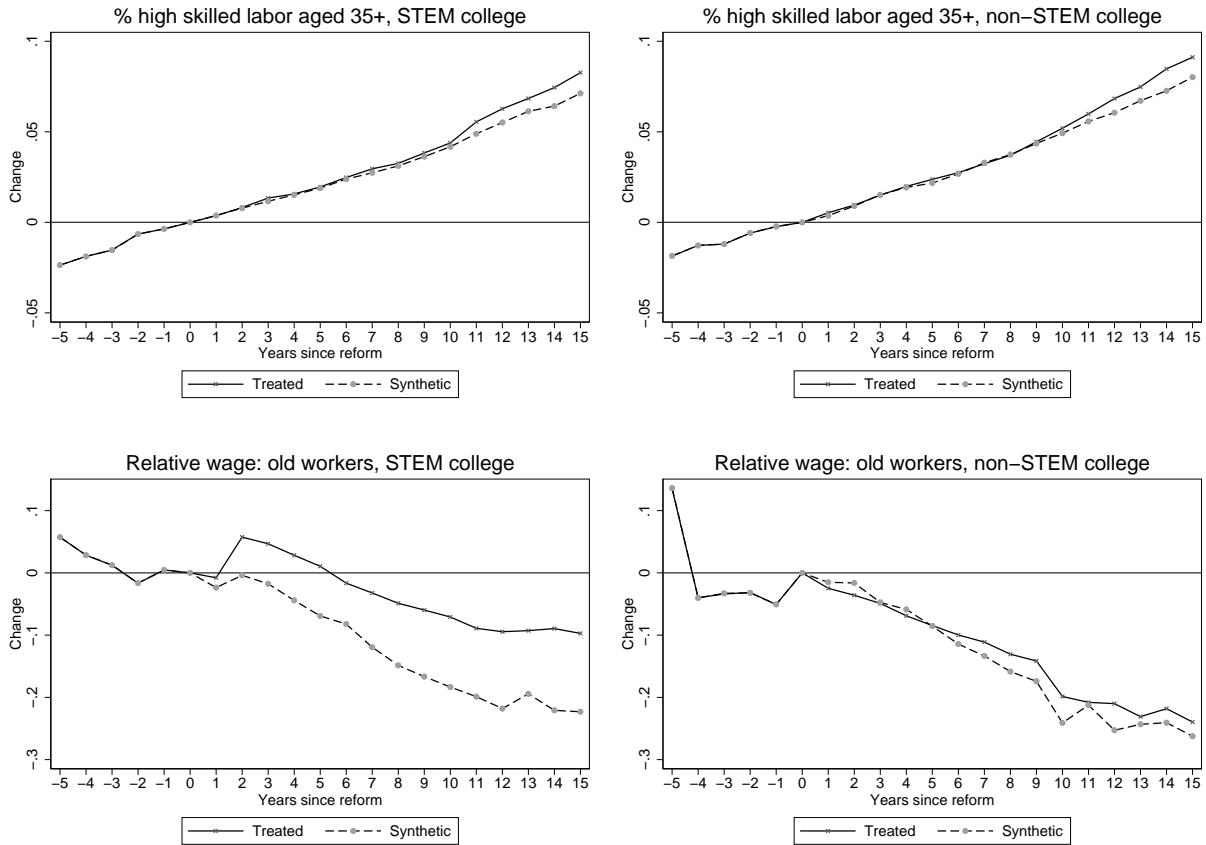
Note: This figure shows the results for the placebo test. The gray lines represent the estimated treatment effect from each placebo of the permutation test. The thick solid line denotes the treatment effect estimated using the actual treated municipalities in the data. The implied p-values of the actual treatment effects are reported in Appendix Table A1. The outcome in the year of the reform is normalized to zero.

Figure 6: The Effects of the Reform on Young Workers: STEM vs Non-STEM Colleges



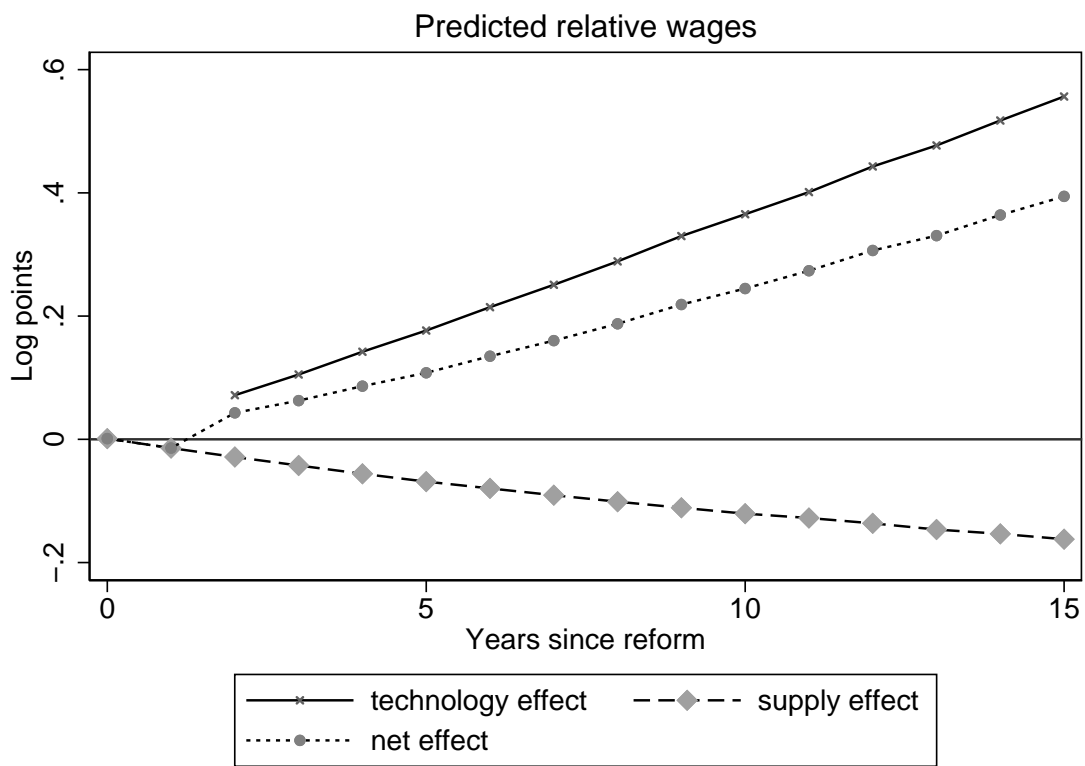
Note: This figure presents the synthetic control estimates on skill composition and wages of young workers, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome variable in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). Each graph reports the weighted average of all treated municipalities and the corresponding synthetic controls, with weight given by the number of plants in the treated municipality in the given year.

Figure 7: The Effects of the Reform on Old Workers: STEM vs Non-STEM Colleges



Note: This figure presents the synthetic control estimates on skill composition and wages of old workers, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome variable in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). Each graph reports the weighted average of all treated municipalities and the corresponding synthetic controls, with weight given by the number of plants in the treated municipality in the given year.

Figure 8: Predicted Relative Wages: Decomposing the Technology and Supply Effects



Note: This figure reports the predicted relative wages from the production function estimates. See Section 5.4 for details.

Table 1: Comparison of Baseline Characteristics Before the Reform, by Treatment Status

	The college reform		
	Treated	Non-treated	Diff. (p-val)
Share of skilled workers	0.118	0.105	0.013* (0.008)
Log high-skilled wages	12.306	12.290	0.016 (0.014)
Log low-skilled wages	11.997	11.995	0.001 (0.015)
Annual change in share of skilled workers	0.003	0.002	0.001** (0.000)
Annual change in log high-skilled wages	-0.001	0.003	-0.004 (0.004)
Annual change in log low-skilled wages	-0.017	-0.015	0.001 (0.002)

Note: This table reports the mean pretreatment characteristics of the municipalities with treatment and the remaining municipalities without treatment. For the top three outcomes, the table reports the weighted averages in 1967. For the bottom three outcomes, the table reports the weighted average annual change between 1967 and 1969. The weights are the number of plants in 1967.

Table 2: Estimates from the Relative Labor Demand Regression

	(1)	(2)	(3)
Aggr. Supply	-0.932** (0.352)	-0.865** (0.368)	-0.690* (0.360)
Trend	0.018 (0.012)	0.016 (0.012)	0.010 (0.012)
Trend×Post	0.015*** (0.004)	0.015*** (0.005)	0.014*** (0.005)
Treated	0.006 (0.018)	-0.024 (0.019)	-0.054** (0.021)
Post	-0.001 (0.015)	0.020 (0.013)	0.035*** (0.012)
Older worker	0.079*** (0.020)	0.079*** (0.021)	0.080*** (0.021)
N	64	64	64

Note: This table reports the estimates of equation (4). See Section 4.3 for details. In column (1), we set $M = 2$ (baseline). In columns (2) and (3), we set $M = 3$ and $M = 4$, respectively.

Table 3: Production Function Estimates

Variable	(1)	(2)	(3)
$\log(S) \times \text{post}$	0.006 (0.031)	0.008 (0.031)	-0.015 (0.026)
$\log(S) \times \text{post} \times \text{trend}$	0.007*** (0.002)	0.007*** (0.002)	0.007** (0.003)
$\log(U) \times \text{post}$	0.013 (0.031)	0.012 (0.032)	0.022 (0.030)
$\log(U) \times \text{post} \times \text{trend}$	-0.009*** (0.003)	-0.009*** (0.003)	-0.008* (0.004)
post	-0.146 (0.301)	-0.150 (0.303)	-0.050 (0.287)
post \times trend	0.038 (0.025)	0.037 (0.025)	0.025 (0.027)
$\log(S)$	0.330*** (0.097)	0.376*** (0.109)	0.230*** (0.083)
$\log(U)$	0.523*** (0.097)	0.476*** (0.111)	0.375*** (0.084)
Log capital ^b	0.163*** (0.020)	0.162*** (0.020)	-
Skill share control function	No	Yes	Yes
LP control function	No	No	Yes

Note: This table reports the estimated production function (see equations (8) and (9) in the text). Variable *Post* is equal to one among treated municipalities in periods at least two years after the reform and zero otherwise. Variable *trend* is the number of years since the reform (normalized to 0 in the year of the reform and set to zero for untreated municipalities). Column (1) reports the estimates without any control function. Column (2) adds the control function for skill compositions and column (3) further controls for unobserved productivity shocks via intermediate inputs. Number of observations = 18441. Standard errors are clustered at municipality level and given in parentheses.

Table 4: Production Function Estimates: Robustness Checks

Variable	(1)	(2)	(3)	(4)
$\log(S) \times \text{post}$	0.014 (0.027)	0.008 (0.031)	0.010 (0.034)	0.018 (0.032)
$\log(S) \times \text{post} \times \text{trend}$	0.005** (0.002)	0.006*** (0.002)	0.007*** (0.002)	0.007*** (0.002)
$\log(U) \times \text{post}$	0.007 (0.033)	0.009 (0.030)	0.003 (0.035)	0.005 (0.068)
$\log(U) \times \text{post} \times \text{trend}$	-0.007** (0.003)	-0.009*** (0.003)	-0.009*** (0.003)	-0.010** (0.004)
post	-0.129 (0.321)	-0.119 (0.260)	-0.046 (0.334)	-0.141 (0.348)
post \times trend	0.026 (0.027)	0.045* (0.025)	0.036 (0.025)	0.033 (0.047)
$\log(K) \times \text{post}$	-	-	-	-0.001 (0.054)
$\log(K) \times \text{post} \times \text{trend}$	-	-	-	0.001 (0.004)

Note: This table reports the estimated production function under alternative specifications (see Section 5.3 in the text). Variable *Post* is equal to one among treated municipalities in periods at least two years after the reform and zero otherwise. Variable *trend* is the number of years since the reform (normalized to 0 in the year of the reform and set to zero for untreated municipalities). Relative to the baseline, Column (1) drops fixed year effects in $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$. Column (2) includes industry fixed effects in $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$. Column (3) adds initial municipality–industry specific characteristics in $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$ (including the share of skilled workers and employment size relative to total manufacturing employment in the municipality, both measured in 1960). Column (4) allows for the reform to impact the productivity of capital. Number of observations = 18441. Standard errors are clustered at municipality level and given in parentheses.

Table 5: College Reform and R&D Activities

	Log total costs of R&D			Log R&D man-years		
	(1)	(2)	(3)	(4)	(5)	(6)
D_{ct}	0.767** (0.362)	0.833*** (0.228)	1.299** (0.552)	0.944*** (0.338)	1.246* (0.677)	1.590 (1.115)
N	859	859	859	469	469	469
Covariates	No S \times Y	Baseline	Munic' trend	No S \times Y	Baseline	Munic' trend

Note: Estimates from equation (10) in the text. Coefficient on D_{ct} shows the effect of the reform. Standard errors are clustered at the municipality level. See Section 6 in the text for details.

APPENDIX

A Implementation of the Synthetic Control Estimation

In this section, we first discuss the estimation of the effect of the reform for a single treated municipality, before aggregating the treatment effects across different municipalities. Following Abadie, Diamond, and Hainmueller (2010), suppose that we observe $J + 1$ municipalities and, without loss of generality, that the first municipality is exposed to the reform from period $T_0 + 1$ on (so the remaining J municipalities are potential controls). Let Y_{ct}^0 be the potential outcome that would be observed for municipality c at time t in the absence of the reform, where $c = 1, \dots, J + 1$, and time periods $t = 1, \dots, T$. Let T_0 be the number of periods before the reform, with $0 < T_0 < T$. Let Y_{ct}^1 be the potential outcome that would be observed for municipality c at time t if the municipality is exposed to the reform from period $T_0 + 1$ to T . Assume that the reform has no effect on the outcome prior to the intervention, so $Y_{ct}^0 = Y_{ct}^1$ for $t \in \{1, \dots, T_0\}$.

Let D_{ct} be an indicator function which takes the value of one if the municipality is treated at time t . The observed outcome is linked to the potential outcomes via

$$Y_{ct} = Y_{ct}^0 + \alpha_{ct} D_{ct} \tag{A.1}$$

where

$$D_{ct} = \begin{cases} 1 & \text{if } c = 1 \quad \text{and} \quad t > T_0; \\ 0 & \text{elsewhere.} \end{cases} \tag{A.2}$$

Let $\alpha_{1t} = Y_{1t}^1 - Y_{1t}^0$ be the effect of the reform for the treated municipality ($c = 1$) at time t in periods $T > T_0$. Because Y_{1t}^1 is observed, in order to estimate α_{1t} , we only need to estimate Y_{1t}^0 . Suppose that Y_{ct}^0 can be parameterized by the following factor model

$$Y_{ct}^0 = \delta_t + \boldsymbol{\theta}_t \mathbf{Z}_c + \boldsymbol{\lambda}_t \boldsymbol{\mu}_c + \varepsilon_{ct} \tag{A.3}$$

where δ_t is an unknown common factor with constant factor loadings across all municipalities, \mathbf{Z}_c is a vector of observed covariates of the municipality that are not affected by the reform, $\boldsymbol{\theta}_t$ is a vector of unknown parameters, $\boldsymbol{\lambda}_t$ is a vector of unobserved common factors, $\boldsymbol{\mu}_c$ is a vector of unknown factor loadings, and the error terms represent unobserved transitory shocks at the municipality level with zero mean. Note that Equation (A.3) generalizes the alternative difference-in-differences model that we also implemented below. The difference-in-differences (fixed-effects) model can be obtained if we impose that λ_t in Equation (A.3) is constant for all t . That is, the difference-in-differences model allows for the presence of unobserved confounders but restricts the effect of those confounders to be constant in time, so they can be eliminated by taking time differences. In contrast, the synthetic control estimator is based on the the factor model, which allows the effects of confounding unobserved characteristics to vary with time.

Abadie, Diamond, and Hainmueller (2010) show that we can use $\sum_{j=2}^{J+1} w_j^* Y_{jt}$ as an estimator for Y_{1t}^0 , where w_j^* is a weight for each potential control municipality j such that

$$\begin{aligned}\sum_{j=2}^{J+1} w_j^* Y_{j\tau} &= Y_{1\tau}, \forall \tau \in \{1, \dots, T_0\} \\ \sum_{j=2}^{J+1} w_j^* \mathbf{Z}_j &= \mathbf{Z}_1 \\ \sum_{j=2}^{J+1} w_j^* &= 1, w_j^* \geq 0\end{aligned}$$

The vector $\mathbf{W}^* = (w_2^*, \dots, w_{J+1}^*)'$ represents a weighed average of the available control municipalities, and therefore, a synthetic control. In practice, it is often the case that no set of weights exists such that these equations hold exactly in the data. Then, the synthetic control observations will be selected so that they hold approximately. Let vector $K = (k_1, \dots, k_{T_0})'$ define a linear combination of pre-intervention outcomes: $\bar{Y}_c^K = \sum_{s=1}^{T_0} k_s Y_{cs}$. For instance, if $k_1 = k_2 = \dots = k_{T_0-1} = 0$ and $k_{T_0} = 1$, then $\bar{Y}_c^K = Y_{cT_0}$. We consider 5 of such linear combinations defined by vectors K_0, \dots, K_4 where $\bar{Y}_c^{K_0} = Y_{cT_0}, \bar{Y}_c^{K_1} = Y_{cT_0-1}, \dots, \bar{Y}_c^{K_4} = Y_{cT_0-4}$. Let $\mathbf{X}_1 = (Z_1'; \bar{Y}_1^{K_0}, \dots, \bar{Y}_1^{K_4})'$ be a $k \times 1$ vector of pre-intervention characteristics for the treated municipality.⁴² Similarly, \mathbf{X}_0 is a $k \times J$ matrix that contains the same variables for the potential control municipalities, where the j th column of \mathbf{X}_0 is $(Z_j'; \bar{Y}_j^{K_0}, \dots, \bar{Y}_j^{K_4})'$. The vector \mathbf{W}^* is chosen to minimize the distance between X_1 and X_0W , where

$$\|\mathbf{X}_1 - \mathbf{X}_0 \mathbf{W}\| = \sqrt{(\mathbf{X}_1 - \mathbf{X}_0 \mathbf{W})' \mathbf{V} (\mathbf{X}_1 - \mathbf{X}_0 \mathbf{W})} \quad (\text{A.4})$$

where V is a symmetric and positive semi-definite matrix. Following Abadie, Diamond, and Hainmueller (2010), we use the “data-driven” V such that the mean squared prediction error of the outcome variable is minimized for the preintervention periods.

The average effect of the reform across R different treated municipalities (i.e. municipalities that ever had a new college in our sampling period) is given by

$$\bar{\alpha}_\tau = \frac{1}{R} \sum_{r=1}^R (\omega(r) \hat{\alpha}_{r,\tau}), \forall T_0 < \tau \leq T \quad (\text{A.5})$$

where $\hat{\alpha}_{r,\tau}$ is the estimated effect of the reform in period τ for municipality r and $\omega(r)$ is the municipality-specific weight that is equal to the number of plants within the municipality.

To assess the extent to which our estimates are statistically important, we follow Abadie, Diamond, and Hainmueller (2010) and estimate a series of placebos by iteratively applying the synthetic control method to every municipality in the pool of potentially control municipalities. Specifically, for each treated municipality, we can perform inference (p-value) at every post-reform

⁴²In vector \mathbf{Z}_j , we include the average share of young workers (aged less than 35) and share of skilled workers, both measured before the time of the intervention.

period as

$$p_{r,\tau} = \frac{1}{J_r} \sum_{j=2}^{J_r+1} \mathbf{1}(\hat{\alpha}_{r,\tau}^{P(j)} < \hat{\alpha}_{r,\tau}), \forall T_0 < \tau \leq T \quad (\text{A.6})$$

where $\hat{\alpha}_{r,\tau}^{P(j)}$ is the effect of the reform when a control municipality j is assigned a placebo reform at the same time as the treated municipality r . $\hat{\alpha}_{r,\tau}^{P(j)}$ is computed following the same procedure outlined for $\hat{\alpha}_{r,\tau}$. Therefore, we can obtain the distribution of placebo effects and the p-value for the treatment effect in municipality r is assessed by computing how the estimated $\hat{\alpha}_{r,\tau}$ ranks in that distribution.

In principle we can calculate and report p-values for the treatment effects of each treated municipality from the empirical distribution of the gaps implied by the placebos. However, it is simpler to present a single p-value for the treatment effects averaged across all treated municipalities.⁴³ We construct a distribution of average placebo effects according to the following steps:

1. For each treated municipality r , we compute all the placebo effects using the available controls corresponding to the municipality r . Each placebo j produces a placebo effect in each period after the reform, denoted by $\hat{\alpha}_{r,\tau}^{P(j)}$.
2. We compute the average placebo effect by randomly selecting (with replacement) a single placebo corresponding to each treated municipality r and then taking the average across the R placebos: $\bar{\alpha}_\tau^k = \frac{1}{R} \sum_{r=1}^R \hat{\alpha}_{r,\tau}^{P(j)}$
3. Repeating step 2 for K times ($K = 50$), we obtain the p-value for the average effect in period τ as $p_\tau = \frac{\sum_{l=1}^K \mathbf{1}(\hat{\alpha}_\tau^l > \hat{\alpha}_\tau^0)}{K}$.⁴⁴

B All Untreated Municipalities as Comparisons: Testing for Pre-treatment Differential Trend

In this section, we report our findings using a standard difference-in-difference research design where all the untreated municipalities are included as comparisons. The identifying assumption is that the geographic location of the college expansion is not correlated with different underlying trends in local labor-market outcomes across municipalities. One way to test this assumption is to check whether the outcome variables in the treated and control municipalities evolve similarly over time during the pre-reform period.

We formalize this design by using a flexible event-study regression specification. This specification allows the analysis to characterize changes in the effect of the reform in the short- and long-run and evaluate the evolution of pre-treatment unobservables in treated municipalities (e.g.,

⁴³Cavallo, Galiani, Noy, and Pantano (2013) uses a similar approach to draw inference using the average placebo treatment effect.

⁴⁴For unskilled wages, we report the p-values using $\frac{\sum_{l=1}^K \mathbf{1}(\hat{\alpha}_\tau^l < \hat{\alpha}_\tau^0)}{K}$. See Appendix Table A1 for details.

Jacobson, LaLonde, and Sullivan (1993)). We estimate the reduced-form impact of the reform using the following linear regression model:

$$Y_{ct} = \theta_c + \gamma_{k(c)t} + \sum_{y=1}^{13} \tau_y D_c \mathbf{1}(t - T_c = y) + \sum_{y=-5}^{-1} \pi_y D_c \mathbf{1}(t - T_c = y) + \varepsilon_{ct} \quad (\text{B.1})$$

where Y_{ct} is the outcome of interest in municipality c ; θ_c is a set of municipality fixed effects; $\gamma_{k(c)t}$ is a set of county-by-year fixed effects, which captures time-varying and county-level changes. The inclusion of county-by-year fixed effects means that we are comparing outcomes across treated and untreated municipalities within the same county. It also implies that the counties with existing universities prior to the reform (such as Oslo) are not used as a comparison group. T_c is the year that college reform is effective in municipality c . The key regressor is the interaction of D_c , a dummy variable equal to one if the municipality ever received a new college, and an indicator function, $\mathbf{1}(t - T_c = y)$, which is equal to one when the year of observation is y years from the reform year T_c .⁴⁵ Standard errors are clustered at the municipality level. The regression is weighted by the number of plants in each municipality.

The event-study specification identifies changes in the effect of the reform over time. The set of parameters π describes the differential evolution of pre-treatment unobservables in treated municipalities (relative to control), that is, the “treatment effects” preceding the reform. They provide an important falsification test on whether any preexisting, unobserved, and nonlinear trends may confound the estimates of the true reform effects τ . For the estimated treatment effects to be internally valid, one requirement of the regression approach is that the estimated π s should be close to zero and statistically insignificant.

In Appendix Tables A3 and A4, we report the effects of the reform on skill shares and log average wages by age groups, respectively, from the regression model as specified in equation (B.1). We allow the effects of the reform (the estimated τ s) to vary by 1–2, 3–4, 5–8, 9–12 and 13 years or more after the reform. Despite the pre-reform differential trend in a few outcomes, the regression estimates are broadly consistent with estimates from the synthetic control estimator. Our results highlight the importance of the appropriate selection of control group that allows us to separate trends from the treatment effects. For instance, among old workers, the estimated increase in skill shares is much less pronounced among the synthetic control estimates than the regression estimates. This is due to the positive preexisting differential trend in the skill shares in the treated municipalities (relative to the control municipalities) prior to the reform. The synthetic control method takes account of this by selecting control municipalities that match this preexisting trend.

⁴⁵The omitted category is the indicator for $y = 0$. Given the period of the panel data and the timing of the reform, to ensure the parameters are well estimated, values of $y < -5$ are grouped to be equal to -5 and all values greater than 13 are grouped into the category 13.

C College Expansion and Cognitive Ability of Young Workers

A challenge to the interpretation of the wage results related to young workers is that changes in the relative number of college graduates may also affect the relative composition of the pool of college graduates. For instance, in response to an opening of a regional college, selection into college education may be based on aptitude to a greater extent than mobility costs. As a result, the “best” of previous upper-secondary graduates may move on to college education, which potentially drives the quality of high-skilled workers upward and the quality of low-skilled workers downward. If the reform leads to changes in unobserved relative quality between the two skill groups in this direction, we may be misinterpreting part of the residual changes in wages to relative demand change (that is, in fact, driven by changes in supply). This concern is also relevant to the interpretation of the estimates being technical change in the production function estimation (discussed in Section 5).

We attempt to address this concern by examining the effect of the reform on cognitive ability, using IQ scores from military draft data that were recorded from several male cohorts upon entering military service.⁴⁶ While the observed IQ score is potentially insufficient to capture individual ability fully, the correlation between the reform and the IQ scores may nevertheless reveal useful information about the importance of composition changes in biasing our productivity estimates. Heckman, Stixrud, and Urzua (2006) have shown that there is a strong correlation between IQ scores and labor-market outcomes.

The IQ scores are available among males starting from the cohort born in 1950. We study the 1950–1964 cohorts who have completed at least some high school education. Individuals in these cohorts are expected to make a college education choice during the period 1969–1983 (at age 19 years). To test whether the IQ scores of cohorts choosing college education after the reform are different from the IQ scores of cohorts whose college education choice was made before the reform, we run the following regression by exploiting variations in the exposure to the reform by cohort and municipality of residence:

$$Q_i = \chi_m + \chi_c + \pi R_{mc} + \epsilon_{imc} \quad (\text{C.1})$$

where Q_i is the IQ score for individual i ; χ_c are fixed cohort effects; χ_m are fixed municipality effects, where municipality is defined as the municipality of residence before college decision (at age 17 years); and R_{mc} is a reform indicator, equal to one if there is a new college established at age 19 years for cohort c residing in municipality m . We estimate equation (C.1) by skill type

⁴⁶Military service is compulsory for all able males in Norway. Before entering the service, their medical and psychological suitability is assessed: this occurs around their 18th birthday. The IQ measure is a composite score from three timed IQ tests: arithmetic, word similarities, and figures. The arithmetic test is quite similar to the arithmetic test in the Wechsler Adult Intelligence Scale (WAIS) (Sundet, Tambs, Harris, Magnus, and Torjussen, 2005; Cronbach, 1964). The word test is similar to the vocabulary test in WAIS, and the figures test is similar to the Raven Progressive Matrix test (Cronbach, 1964) (also see Sundet, Barlaug, and Torjussen (2004), Sundet, Tambs, Harris, Magnus, and Torjussen (2005), and Thrane (1977) for details). The composite IQ test score is an unweighted mean of the three sub-tests. The IQ score is reported in stanine (Standard Nine) units, a method of standardizing raw scores into a nine-point standard scale with a normal distribution, a mean of five, and a standard deviation of two.

of the individual, and the estimated coefficient π^g is the effect of the reform on IQ scores among individuals with skill level g .

Appendix Table A5 presents the estimated coefficients π for different outcome variables. Column (1) shows that the reform increases college attainment, consistent with the hypothesis that certain upper-secondary graduates move on to college education as a result of the reform. More important for our purpose, columns (2)–(3) do not report any significant evidence that the reform changes the average IQ scores among each of the skill groups. Taken together, although the reform shifts marginal students from upper-secondary education to a college education, there is no evidence that the average quality of graduates changes after the reform. Mobility cost appears to be a more important factor than cognitive ability in the selection of students into these regional colleges.

D Construction of Skill Shares Within Plants

Given data limitations, discussed in Section 2.2, we do not observe skill shares every year at the industry–municipality level, π_{jct} . We only observe this variable every 10 years in the Census. However, from the register data, we have annual data on skill shares at the level of the municipality, π_{ct} . The relationship between the skill share at industry–municipality level and the skill share at municipality level can be written as:

$$\pi_{jct} \equiv \frac{H_{jct}}{T_{jct}} = \left(\frac{H_{jct}}{H_{ct}}\right) \times \left(\frac{T_{ct}}{T_{jct}}\right) \times \left(\frac{H_{ct}}{T_{ct}}\right) \equiv Q_{jct} \times \left(\frac{H_{ct}}{T_{ct}}\right) \equiv Q_{jct} \times \pi_{ct} \quad (\text{D.1})$$

where H_{jct} and T_{jct} are total skilled labor and total employment in industry j , municipality c , and year t , respectively, and H_{ct} and T_{ct} are total skilled labor and total employment in municipality c and year t , respectively.

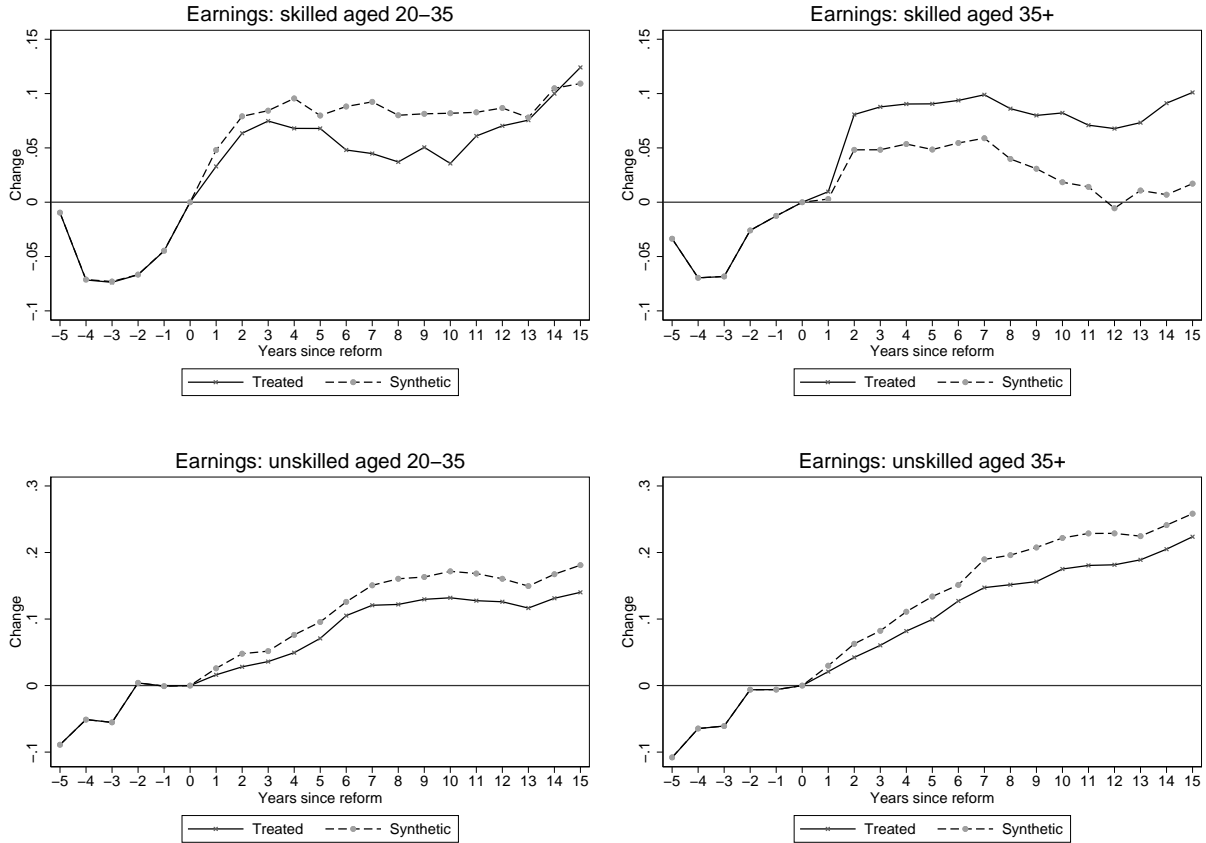
Using decennial data on π_{jct} combined with annual data on π_{ct} , we construct annual observations of π_{jct} using the following steps:

Step 1: In each year of the decennial data, we compute π_{jct} for each municipality and industry (at ISIC 3-digit level) and π_{ct} for each municipality. We then calculate Q_{jct} , where Q_{jct} can be interpreted as the rate of pass-through from a change in municipality-level skill share to the skill share of industry j within the same municipality.

Step 2: By fixing Q_{jct} at its level in 1960 (a pre-reform year) within each industry and municipality, we predict the π_{jct} in each year other than those contained in the decennial data. Therefore, change in skill shares for an industry within a given municipality is a constant proportion of change in skill shares in that municipality.

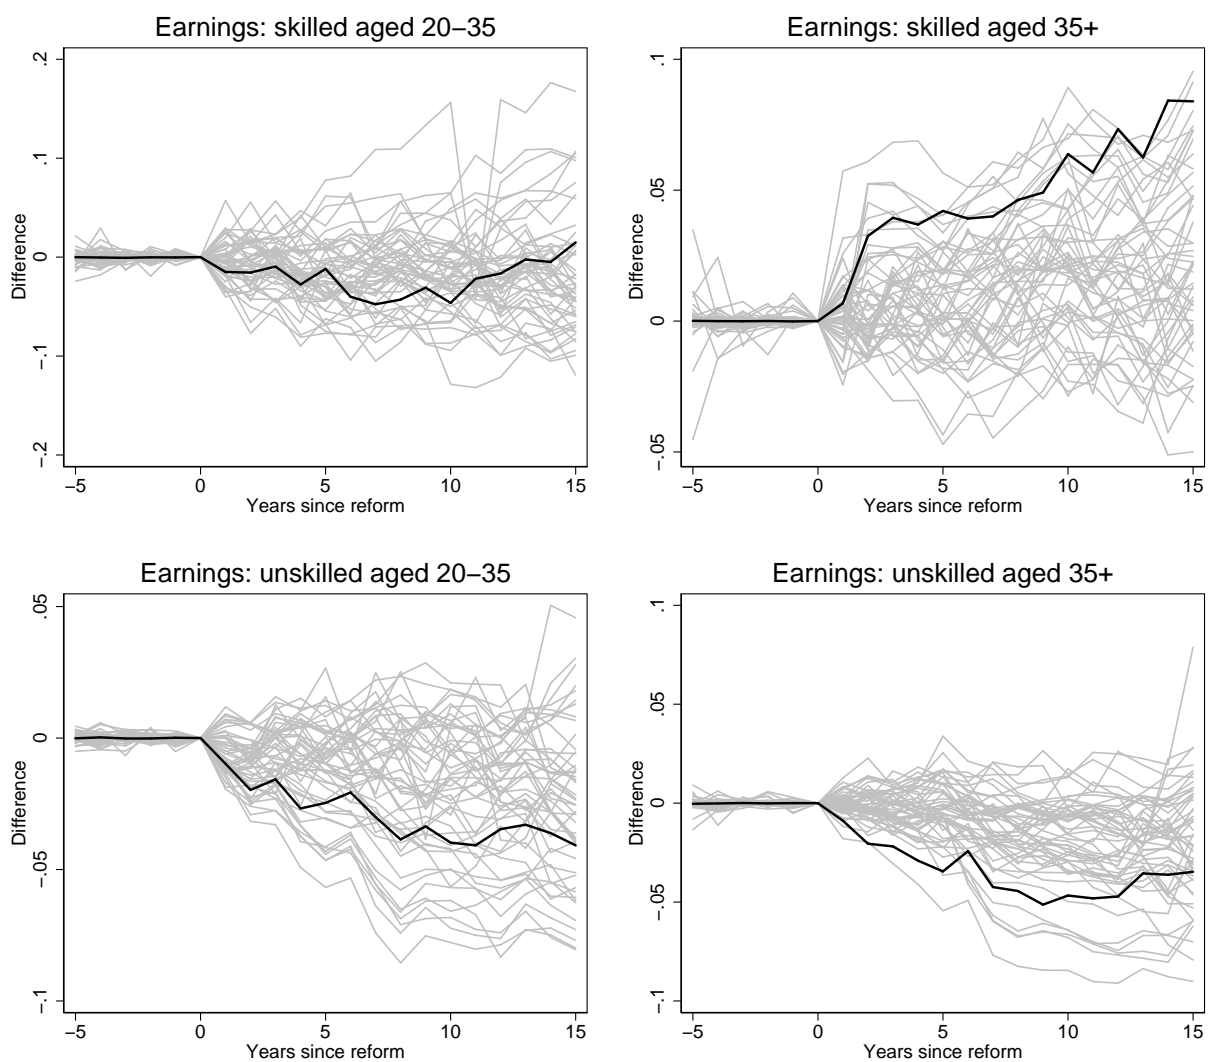
APPENDIX TABLES AND FIGURES

Figure A1: The Effects of the Reform on Absolute Wages and Skill Compositions



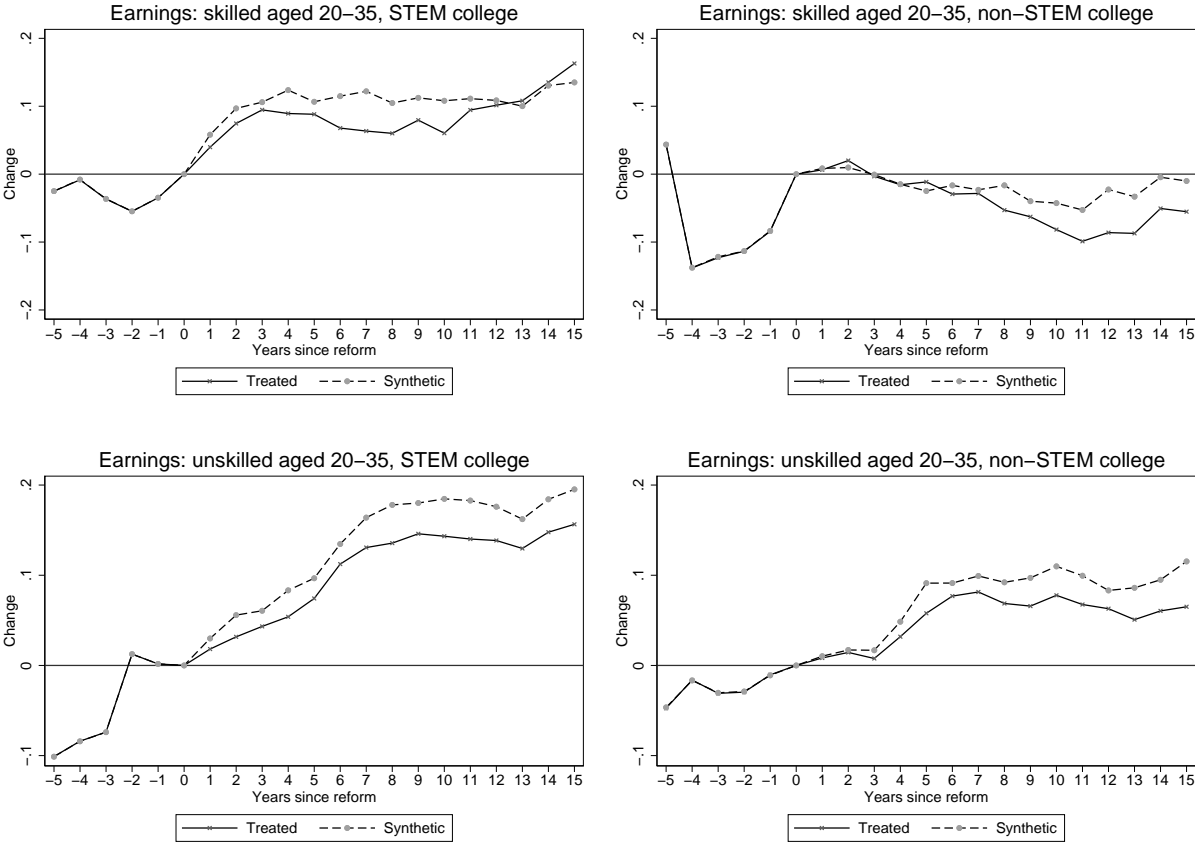
Note: This figure presents the synthetic control estimates on skill composition and absolute wages of the workforce. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). Each graph reports the weighted average of all treated municipalities and the corresponding synthetic controls, with weight given by the number of plants in the treated municipality.

Figure A2: The Effects of the Reform on Absolute Wages and Skill Compositions: Placebo Tests



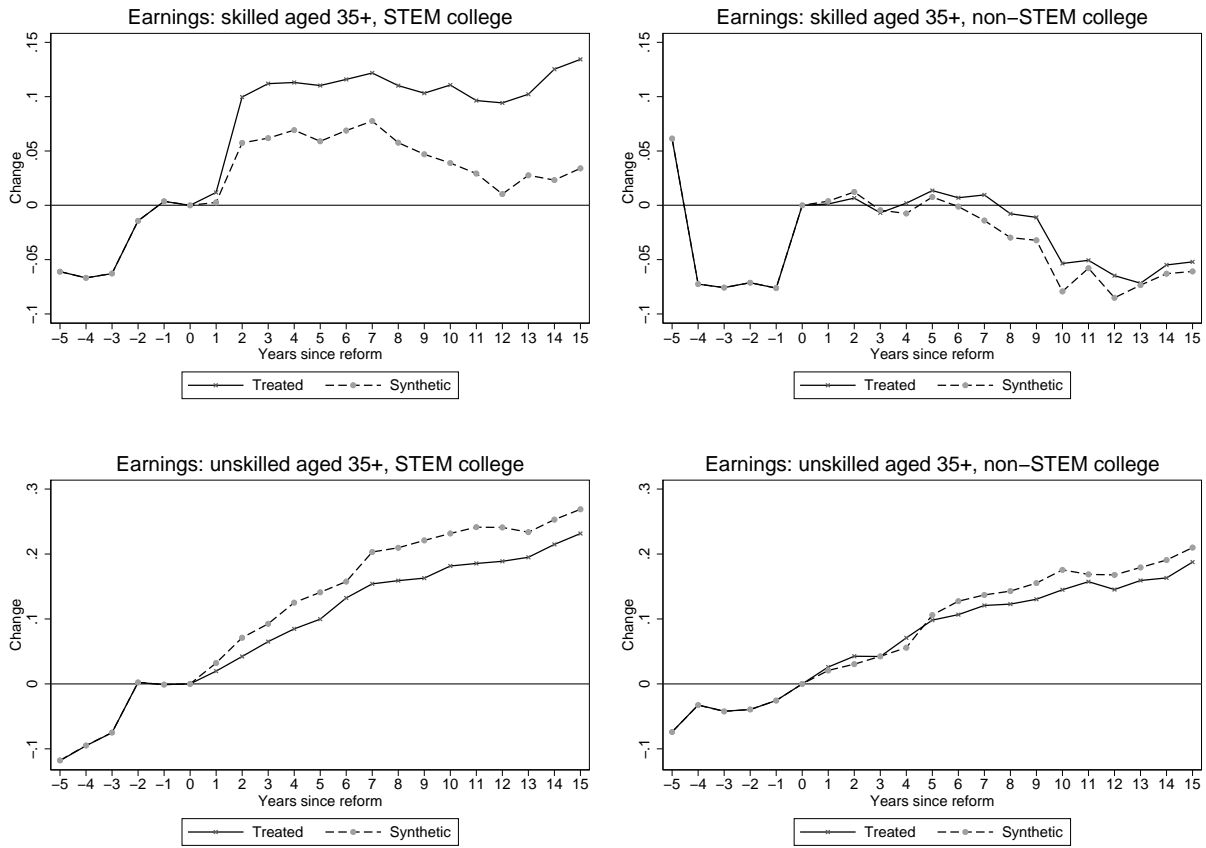
Note: This figure shows the results for the placebo test. The gray lines represent the estimated treatment effect from each placebo of the permutation test. The thick solid line denotes the treatment effect estimated using the actual treated municipalities in the data. The implied p-values of the actual treatment effects are reported in Appendix Table A1. Outcome in the year of the reform is normalized to zero.

Figure A3: The Effects of the Reform on Absolute Wages of Young Workers: STEM vs Non-STEM Colleges



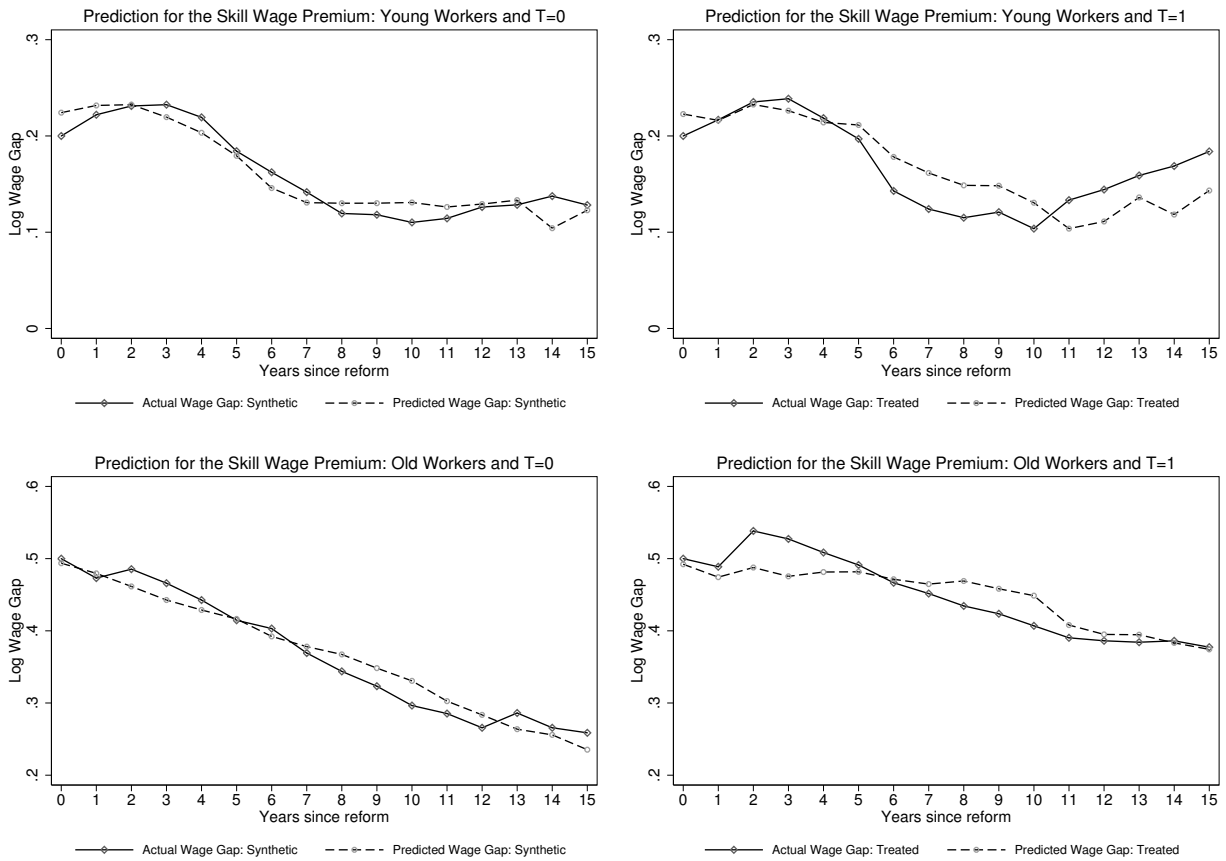
Note: This figure presents the synthetic control estimates on skill composition and wages of young workers, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome variable in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). Each graph reports the weighted average of all treated municipalities and the corresponding synthetic controls, with weight given by the number of plants in the treated municipality in the given year.

Figure A4: The Effects of the Reform on Absolute Wages of Old Workers: STEM vs Non-STEM Colleges



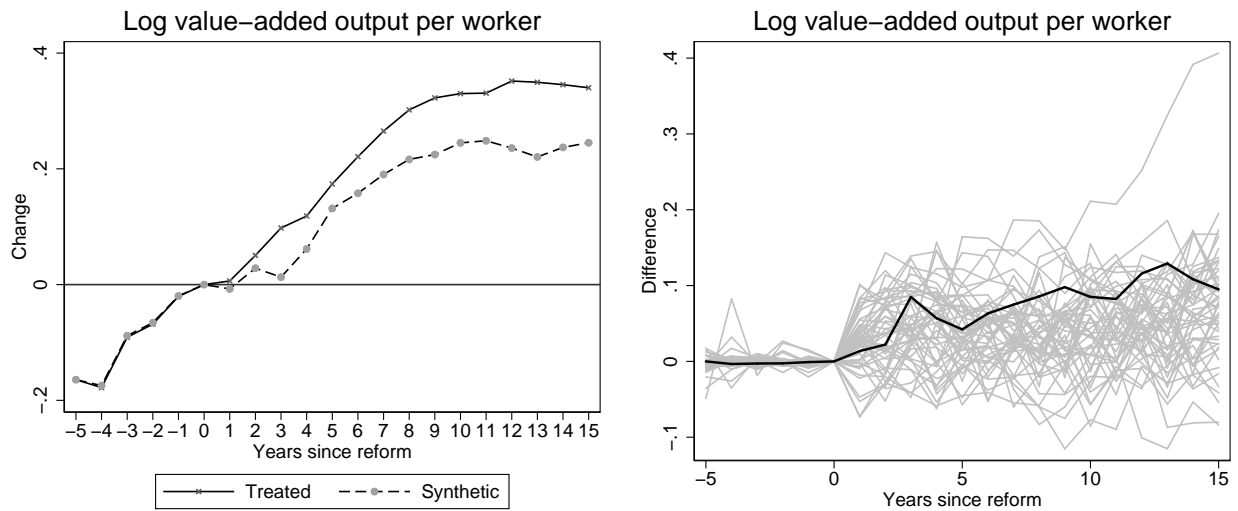
Note: This figure presents the synthetic control estimates on skill composition and wages of old workers, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome variable in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). Each graph reports the weighted average of all treated municipalities and the corresponding synthetic controls, with weight given by the number of plants in the treated municipality in the given year.

Figure A5: Prediction for the Skill Wage Premium



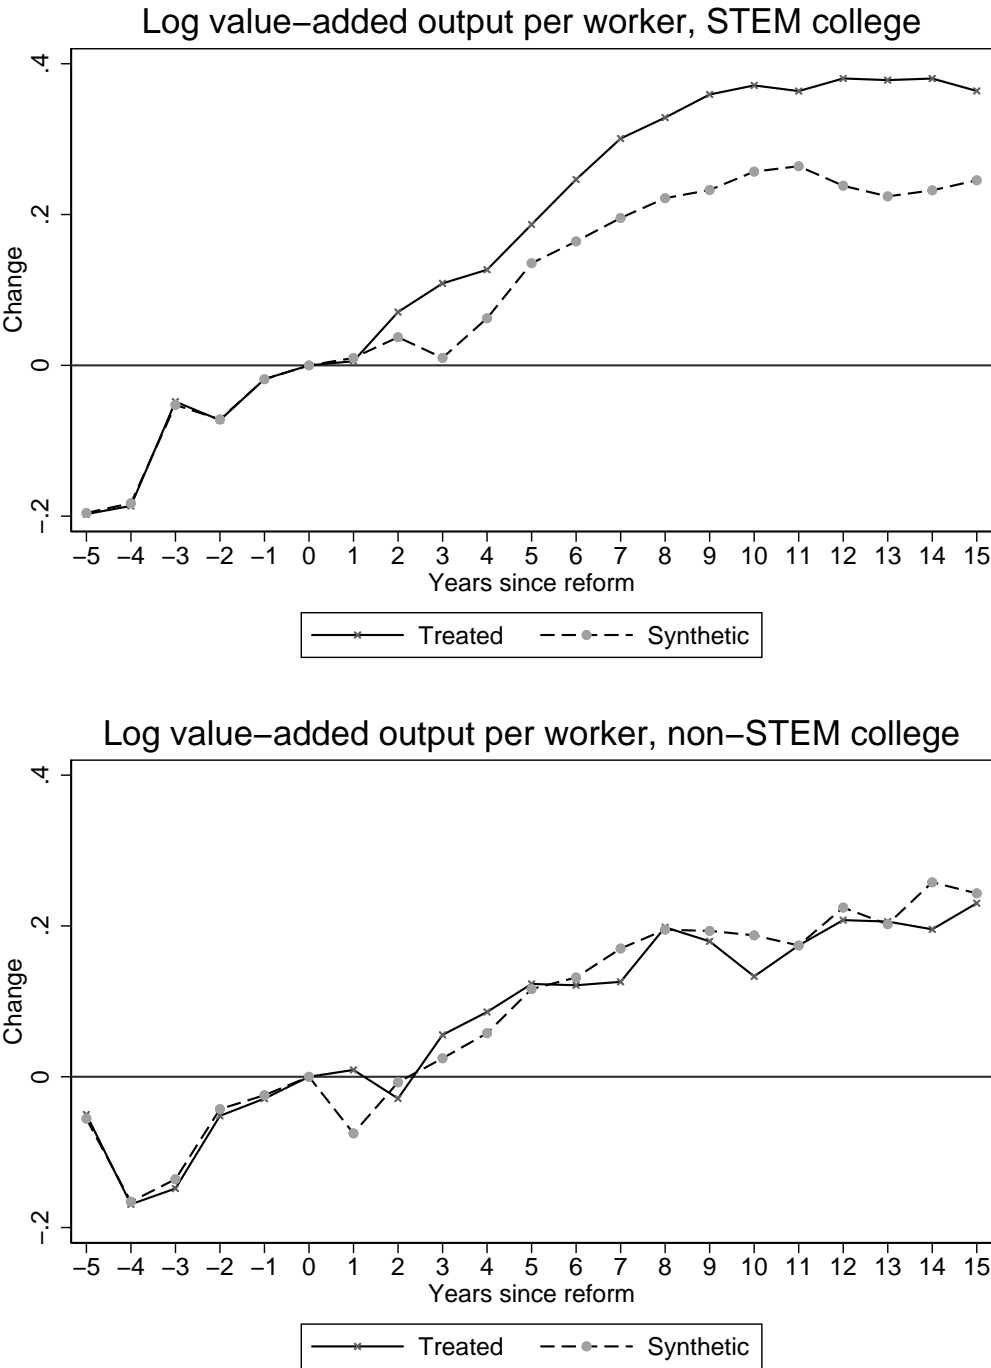
Note: This figure presents the fitness of the relative demand model (equation 4) to the skill premiums of the treated group and synthetic control group estimated by the synthetic control analysis. See Section 4.3 for details.

Figure A6: The Effects of the Reform on Value-added Output per Worker



Note: The graph to the left presents the synthetic control estimates on log value-added output per worker. Each data point represents the mean outcome in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). The graph to the right shows the results for the placebo test. The gray lines represent the estimated treatment effect from each placebo of the permutation test. The thick solid line denotes the treatment effect estimated using the actual treated municipalities in the data. The implied p-values of the actual treatment effects are reported in Appendix Table A1. On both graphs, the year of the reform is normalized to period zero.

Figure A7: The Effects of the Reform on Value-added Output per Worker: STEM vs Non-STEM Colleges



Note: This figure presents the synthetic control estimates on log value-added output per worker, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). Each graph reports the weighted average of all treated municipalities and the corresponding synthetic controls, with weight given by the number of plants in the treated municipality.

Table A1: Implied p-values from the Permutation Tests

# of years post-reform	Share of skilled workers		Relative wages			Absolute wages			Value-added per worker
	Young	Old	Young	Old	S, young	S, old	U, young	U, old	
1	.28	.34	.54	.36	.74	.32	.42	.06	.66
2	.12	.2	.26	.08	.6	.14	.16	0	.56
3	.1	.06	.24	.1	.54	.1	.28	.02	.28
4	.12	.26	.24	.06	.64	.1	.18	.02	.5
5	.16	.24	.16	.04	.44	.02	.22	.04	.42
6	.16	.3	.54	.08	.78	.14	.28	.14	.26
7	.12	.24	.58	.06	.8	.14	.26	.1	.18
8	.04	.28	.56	.08	.8	.12	.22	.1	.16
9	.06	.28	.34	.06	.66	.16	.32	.1	.12
10	.06	.24	.44	.06	.82	.08	.26	.12	.34
11	0	.12	.26	.1	.5	.12	.24	.12	.26
12	0	.1	.24	.1	.38	.02	.28	.14	.1
13	0	.08	.16	.1	.38	.06	.28	.2	.06
14	0	.02	.14	.08	.34	0	.3	.26	.24
15	0	0	.04	.1	.2	.04	.26	.3	.36

Note: This table reports the p-values implied by the permutation test of the synthetic control estimator. See Appendix Section A for details. The p-value in each period τ is calculated by $\frac{\sum_{i=1}^K \mathbf{1}(\hat{\alpha}_\tau^i > \hat{\alpha}_\tau^0)}{K}$, where K is the number of placebos (we use $K = 50$); $\hat{\alpha}_\tau^i$ is the treatment effect from each placebo; and $\hat{\alpha}_\tau^0$ is the estimated treatment effect (reported in Figure 5). For unskilled wages, we report the p-values using $\frac{\sum_{i=1}^K \mathbf{1}(\hat{\alpha}_\tau^i < \hat{\alpha}_\tau^0)}{K}$.

Table A2: Local Labor Demand Estimation: First-step Estimates

	(1)
Age-group specific relative supply	-0.591*** (0.062)
Year effects:	
Year 1	0.011 (0.021)
Year 2	0.042* (0.021)
Year 3	0.052** (0.022)
Year 4	0.053** (0.022)
Year 5	0.064*** (0.022)
Year 6	0.046** (0.022)
Year 7	0.064*** (0.022)
Year 8	0.084*** (0.022)
Year 9	0.099*** (0.022)
Year 10	0.113*** (0.022)
Year 11	0.143*** (0.023)
Year 12	0.157*** (0.023)
Year 13	0.148*** (0.023)
Year 14	0.166*** (0.023)
Year 15	0.182*** (0.024)
N	32

Note: This table reports the estimates using equation (3) in the text.

Table A3: Regression Estimates: Young Workers

	Share of skilled workers (1)	Relative wage (2)	Log skilled wage (3)	Log unskilled wage (4)
Post-reform				
Years 1 to 2	0.003 (0.002)	0.004 (0.019)	-0.004 (0.018)	-0.008** (0.004)
Years 3 to 4	0.006** (0.003)	0.012 (0.016)	-0.002 (0.017)	-0.014*** (0.004)
Years 5 to 8	0.013*** (0.004)	0.005 (0.016)	-0.013 (0.016)	-0.017** (0.008)
Years 9 to 12	0.027*** (0.006)	0.001 (0.014)	-0.016 (0.017)	-0.016 (0.011)
Years 13+	0.037*** (0.008)	0.039* (0.021)	0.020 (0.027)	-0.019 (0.012)
Pre-reform				
Years 2 to 1	-0.005** (0.002)	0.001 (0.016)	0.002 (0.014)	0.000 (0.005)
Years 3 to 4	-0.006* (0.003)	0.000 (0.024)	0.003 (0.018)	0.002 (0.011)
Years 5 or before	0.008** (0.003)	0.042** (0.019)	0.030** (0.014)	-0.012 (0.016)
N	9127	9127	9127	9127

Note: This table reports estimates using equation (B.1) in the text. The unit of observation is a municipality-year. The regression is weighted by the number of plants in the municipality. Standard errors are clustered at the municipality level.

Table A4: Regression Estimates: Old Workers

	Share of skilled workers (1)	Relative wage (2)	Log skilled wage (3)	Log unskilled wage (4)
Post-reform				
Years 1 to 2	0.002*** (0.001)	0.006 (0.006)	0.004 (0.004)	-0.003 (0.003)
Years 3 to 4	0.004*** (0.001)	0.019** (0.008)	0.012** (0.006)	-0.007 (0.005)
Years 5 to 8	0.006*** (0.002)	0.031*** (0.008)	0.015*** (0.006)	-0.016*** (0.005)
Years 9 to 12	0.012*** (0.003)	0.039*** (0.011)	0.023** (0.009)	-0.016** (0.008)
Years 13+	0.025*** (0.003)	0.041*** (0.012)	0.026* (0.014)	-0.015 (0.013)
Pre-reform				
Years 2 to 1	-0.001 (0.001)	-0.012 (0.008)	-0.010 (0.006)	0.001 (0.004)
Years 3 to 4	-0.008*** (0.003)	-0.017 (0.017)	-0.016 (0.011)	0.001 (0.010)
Years 5 or before	-0.000 (0.005)	0.021 (0.018)	0.008 (0.008)	-0.013 (0.014)
N	9127	9119	9119	9127

Note: This table reports estimates using equation (B.1) in the text. The unit of observation is a municipality-year. The regression is weighted by the number of plants in the municipality. Standard errors are clustered at the municipality level.

Table A5: College Reform, College Education, and Average IQ Scores by Skill Groups

	College education (1)	IQ scores conditional on College (2)	High school (3)
Reform	0.033*** (0.009)	0.051 (0.082)	0.030 (0.043)
N	293488	111761	161249

Note: Estimates using equation (C.1) in the text. The sample consists of male individuals born between 1950 and 1964 who have completed at least some high-school education. Standard errors are clustered at the municipality level.

Table A6: Production Function Estimates: Exogenous Supply

Variable	(1)	(2)
$\log(S) \times 0-1$ years post-reform	0.008 (0.024)	-0.011 (0.022)
$\log(S) \times \text{post}$	0.010 (0.034)	-0.018 (0.031)
$\log(S) \times \text{post} \times \text{trend}$	0.007*** (0.002)	0.007** (0.003)
$\log(U) \times 0-1$ years post-reform	-0.034 (0.028)	-0.012 (0.027)
$\log(U) \times \text{post}$	0.003 (0.035)	0.019 (0.037)
$\log(U) \times \text{post} \times \text{trend}$	-0.009*** (0.003)	-0.008* (0.004)
0-1 years post-reform	0.375 (0.257)	0.262 (0.256)
post	-0.046 (0.334)	0.022 (0.330)
post \times trend	0.036 (0.025)	0.023 (0.028)
LP control function	No	Yes

Note: This table reports the estimated production function without the control function for skill composition. Variable *Post* is equal to one among treated municipalities in periods at least two years after the reform and zero otherwise. Variable *trend* is the number of years since the reform (normalized to 0 in the year of the reform and set to zero for untreated municipalities). In both columns, we include a dummy for 0-1 years after the reform and interact it with skilled- and unskilled-labor input. Column (2) includes the control function for unobserved productivity shocks via intermediate inputs. Number of observations = 18441. Standard errors are clustered at the municipality level and given in parentheses.