THE SUPPLY OF SKILL AND ENDOGENOUS TECHNICAL CHANGE: EVIDENCE FROM A COLLEGE EXPANSION REFORM

Pedro Carneiro
University College London,
Institute for Fiscal Studies,
Centre for Microdata Methods and Practice

Kai Liu
Faculty of Economics,
University of Cambridge,
Department of Economics,
Norwegian School of Economics

Kjell G. Salvanes
Norwegian School of Economics,
CELE and FAIR NHH

We examine the labor market consequences of an exogenous increase in the supply of skilled labor in several municipalities 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. (JEL: J23; J24; O33; I24.)

Acknowledgments: We thank Daron Acemoglu, 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. We thank Hans Hvide for sharing the patent data set. 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 through the Centre for Microdata Methods and Practice (ES/P008909/1), and support of the European Research Council through ERC-2015-CoG-682349. Kjell G. Salvanes also gratefully acknowledges the support by the Research Council of Norway Grant 202445.

E-mail: p.carneiro@ucl.ac.uk (Carneiro); kai.liu@econ.cam.ac.uk. (Liu); kjell.salvanes@nhh.no (Salvanes)

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 the inflow of skill 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.

We document three main empirical results. 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.

These empirical results are 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. Moreover, 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.[1]

1. The Card and Lemieux (2001) model is stylized: it decomposes relative wage changes into contributions from relative supply changes and relative demand changes, where the relative demand changes are labelled as “technology changes". It is agnostic with regard to the specific channel driving the

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 production functions, relying on information on output and input factors (and 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 (we document that no new patents are registered). 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. Moreover, we also show that the construction of new colleges is unlikely to have caused substantial increases in the demand of skilled individuals since the employment in the college sector was tiny relative to the size of the labour markets in the locations where new colleges were established. Finally, we explain that our findings are not affected by migration since migration does not respond to college construction, and we are also able to rule out any trade based explanations of changes in factor prices.

Our empirical analysis is based on several population-wide and long panel data sets, containing rich firm-level information on inputs and outputs, and individual-level information on demographics and labor market outcomes. Firm-level data is available for the population of plants in the manufacturing sector in Norway, spanning the years between 1967 and 1990. Individual-level data combine several administrative registers covering all adult individuals in Norway from the same period. We use the individual level data to construct time-series of wage and labor supply by skill groups for each municipality and time period (we consider each municipality to be a different 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 experiencing a college opening over the period we study than not experiencing a college opening, and this method enables us to find appropriate control municipalities for each municipality with a college opening. Our main results are robust to using a standard difference-in-difference estimator. We model the demand of

technical change. More specifically, the technical change identified from the Card and Lemieux framework cannot reveal whether the technical change is driven by any human capital spillover effects or firms using technology that is more complementary to skilled labor. Our firm-level analysis provides more direct evidence on the technology channel, and we discuss the human capital spillover literature at end of this section.

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.[2] 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 unbundle the influence of these two forces on skill prices, 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. In our benchmark model, 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.

Our findings are robust to changes in the how long the delayed response to SBTC is assumed to be. In addition, we also generate similar results from a model where we do not estimate the elasticity of substitution between skilled and unskilled workers, and instead use reasonable estimates of this parameter from the literature. This allows us to avoid having to assume any delay in the SBTC response when estimating the impact of SBTC on the demand for skilled workers.

The impact of the reform on firm-level productivity is estimated using standard production function techniques, and the impact on firms' R&D activities is studied using a typical difference-in-difference estimator. This is because our estimation of the structural parameters of firm-level production functions is not as amenable to a synthetic control estimator as the estimation of reduced-form labor market impacts of the reform. In turn, our R&D data do not allow us to use a sufficient number of pre-reform years needed for a credible implementation of the synthetic control estimator.

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); Beaudry and Green (2003); Blundell, Green, and Jin (2018)). We provide new evidence that endogenous SBTC responses to shocks in the supply of skill led to quantitatively large increases in the skill premium, consistent with a strong relative equilibrium bias (Acemoglu, 2007). In the long run, our results are also consistent with a strong absolute equilibrium bias where the marginal products of skilled workers increase and the demand curves for skilled labor become upward-sloping. Our work is complementary to recent empirical papers on the reaction of

2. We also explore the implications of other models of local labour markets, including a simple multi-sectoral trade model. We find them to be inconsistent with our empirical findings.

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 and our evidence points to endogenous technical change for a large group of workers, which is more general that the existing studies (which tend to focus on specific industries). Our work also 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 analyzing local economy adjustments to labor supply shocks. That literature often uses immigration flows in a local labor market as a source of 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).[3] 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, consequently, on the wages of skilled workers.

Our paper also closely relates to a growing literature estimating the spillover effects of education.[4] Most of the earlier papers in this literature attempted to estimate the size of spillovers from education by comparing the wages of otherwise similar individuals who work in cities or states with different average levels of education (e.g. Rauch (1993); Acemoglu and Angrist (2001); Ciccone and Peri (2006); Moretti (2004a)). One exception is Moretti (2004b), who estimates education spillover effects on the productivity of manufacturing plants in the US. He finds that productivity of

3. A 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.

4. In the macro-growth literature, recent papers have provided empirical evidence pointing to the large role of human capital externality in explaining regional variation in development (e.g., Acemoglu and Dell (2010).

plants in cities that experience large increases in the share of college graduates rises more than the productivity of similar plants in cities that experience small increases in the share of college graduates. By combining the experimental variation from the college expansion reform and the production function estimation, we extend Moretti (2004b) to estimate the effects of education spillovers on both factor-neutral *and* factor-biased productivity. The nature of our panel also allows us to examine the dynamic changes in productivity both in the short- and the long-run.[5] The spillover effects of education in our paper work through endogenous technology investment responses by firms, which is a possible (but not the most standard) interpretation for education spillovers.

More recently, Kantor and Whalley (2014, 2019) study the spillover effects of research innovations from higher education sector. Kantor and Whalley (2014) find that an exogenous increase in an university's spending increases the local average wage in the non-education sector, particularly among industries which employ more college graduates. Kantor and Whalley (2019) find that proximity to the newly established agricultural experiment stations at land-grant U.S. colleges in the late 19th century had large positive effects on land productivity for two decades. They show that knowledge spillovers increase with literacy rate, consistent with the hypothesis that the knowledge spillovers are skill-biased.[6] Our paper complement these papers by providing new evidence on how an exogenous increase in supply of skills can affect technology change and local labor market. Our results are driven by an increase in the supply of skilled workers, rather than any research innovations produced by the new colleges (which was the focus in the two aforementioned papers). One limitation in our study is that we do not observe the timing of adoption of a specific new technology at the firm level. Although all the evidence from our worker- and firm-level data and R&D data point to endogenous technical change as the most plausible interpretation, we cannot completely rule out the alternative interpretations of human capital spillovers.

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. Section 3 presents evidence on the supply of skill and wages at the local labor market level and interprets the results using the model of Card and Lemieux (2001). In Section 4, we provide plant-level evidence on endogenous technical change via estimating production functions. Section 5 presents further evidence that college openings induced firms to invest more in R&D activities. The last section concludes.

5. In Moretti (2004b), the panel of plants consists of observation in two years (in 1982 and 1992).

6. Davis and Dingel (2019) provide a theoretical model that could generate human capital externalities that are "skill-biased". In their model, individuals can choose to produce either tradable or non-tradable (service) goods. Crucial to the model are the assumptions where (i) tradable productivity is positively correlated with individual's ability and participation in idea exchanges and (ii) individual ability and local learning opportunities are complements in the production function for tradable goods. These assumptions imply that the output gain from greater ability is increasing in local learning opportunities, generating human capital externalities that are "skill-biased".

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. 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)).[7]

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 the country did not have sufficient resources to build new universities. In 1966, a committee appointed by the Government (the Ottosen Committee) proposed to expand the higher education sector by opening regional colleges aiming to provide shorter (two and three years) college education programs. In its report, the Ottosen committee proposed to divide the country into 12 educational regions. Four of the regions already had universities and the committee proposed that the remaining eight regions should each have one regional college (Ottosen-committee, 1966-1970).[8]

Following proposals from the Ottosen Committee, in 1968, the Parliament voted for the opening of four new regional colleges.[9] For the first batch of new colleges, the Parliament initially agreed to an experimental period of five years, followed by an evaluation. However, in 1970, the Parliament decided to expand the reform to two more regions[10] and, through the 1970s, the establishment of regional colleges was expanded to all the educational regions.

The report from the Ottosen Committee also suggested three criteria governing the geographic location of a new college within each educational region. First, new

7. Appendix Figure A1 shows that the overall educational level increased dramatically in Norway from 1960 to 1990. In 1960, only 4.2% of the population above the age of 16 had a college or university degree and 16.3% had a high-school diploma, so the ratio of college-educated persons to high school-educated persons was about 0.26. 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 relative 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).

8. In 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. This is the main procedure determining the total amount of resources allocated to the higher education sector and resource allocations to each college.

9. Three 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).

10. These 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)

colleges should be geographically dispersed across the country. Appendix Figure A2 shows the geographic location of the new colleges across the country.[11] This criterion is clearly met for all the new colleges. Second, new colleges had to be established in regions and in municipalities where all the necessary infrastructure could be completed within a year of the establishment decision. Since this was a necessary requirement to get started within the given time limit, it is clear that this criterium was also met. Third, new colleges should be placed to stimulate growth in regions with stagnation problems. We do not find much empirical support for this last criteria (see Table 1 where we show the mean of sector compositions, manufacturing output and labor market outcomes in the non-treated areas, together with the difference between the treated areas and non-treated areas). Our interpretation is that the two first criteria dominated the selection for placement of the colleges.

In Appendix Section A, we provide additional details on the characteristics of the new colleges, in terms of their sizes and academic programs. We also show that there was very little research output produced by these regional colleges in the period under study.

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.[12]

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.[13] In addition 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

11. Note 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.

12. The earliest year of our plant-level and worker-level data begins in 1967. It could be 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. For this reason, we limit our sample period to 1990. During this period, there were no major reforms to the higher education sector.

13. We explain the use of these data in Appendix Section D.

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.[14] Annual labor earnings are the sum of pre-tax 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. 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. Top-coding is only performed at very high earnings levels, which is another advantage over other similar data sets such as the social security data in the US. However, the top-coded amount increased in the 1970s and early 1980s. To make sure that our average earnings series are comparable over time, we manually top code the earnings data for all earnings above the 96th percentile of the earnings distribution by year.[15]

Education attainment is reported by the education authorities 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 measure the highest level of completed education for each individual in each year.[16]

The sample of individuals being analyzed includes workers aged 20–62 years and whose annual earnings are above the basic amount (2G) required to participate in the national social insurance program.[17] In each year, we classify workers into two

14. Similar 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.

15. Top coding exists in all years prior to 1986, except for 1981. The 96th percentile of earnings distribution by year are not subject to top coding across all years, which we use as our top coding thresholds. For earnings above the top coding threshold, the top coded amount we use is equal to the earnings at the 96th percentile in each year times 1.32, where 1.32 is obtained by dividing the average earnings above the top-coding threshold in 1981, by the value of that threshold (there is no top coding in 1981). In the previous version of the paper, we only trim the top 0.1% observations with the highest earnings, by year. The top coding amount (i.e., the censoring point) went up starting 1971, which leads to a discrete jump in the mean earnings and the jump is bigger for skilled-earnings than unskilled-earnings in that year (because a higher fraction of skilled workers is subject to top-coding prior to 1971). Because year 1971 corresponds to the second year post reform for 5 new colleges, the previous version of the paper documented a discrete jump in mean and relative earnings in both the treated and synthetic group in year 2 (e.g. see Figure A1), which is in fact driven by changes in top-coding in the data.

16. The 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.

17. Although 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, 2017). 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. We validated our sample of workers by linking our

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.[18] By combining workers' information on skill levels and municipality of residence, we construct measures of skill composition and wages in a local labor market over time. 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.[19] The response rate is extremely high because firms are required by law to submit their survey responses.[20] 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, completing at least a total of 5000 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.[21] 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 per firm and year, defined as the gross value of production minus the costs of materials and services.[22]

earnings data with 1970 census, which contains categorical information on annual hours of work. More than 96% of individuals in our data are full-year workers.

18. Our 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)).

19. An 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 Foyn (1991) for a detailed description.

20. The 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.

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

22. Value-added 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).

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 much variation over time in this variable. 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 the capital stock in subsequent years by adding up the value of net investment (in buildings and equipment) in each year and assuming that the current equipment and buildings depreciate at a constant rate.[23] 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.

To measure labor input, we use the total hours of employment for each plant. Unfortunately, from the firm-level data we cannot 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, since 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 4.

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. The R&D sample includes mainly manufacturing firms above a certain size-class.[24] We have access to 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 5, 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 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

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

24. The 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 Møen (2005).

three additional colleges that were built in the same period but were not part of the recommendation from the Ottosen Committee (and are similar to the colleges originally mentioned in the report). The first college opening occurred in year 1967 and the last college opened in 1981.

3. 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 relative wages of skilled workers in local labor markets. We interpret our estimates using the model of Card and Lemieux (2001). Our results suggest that for several years following the establishment of a college there is significant skill-biased technology change in the affected labor markets. In Section 4, 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.

3.1. Construction of the Control Group

There are only 15 municipalities that benefited from a college opening during the reform period 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 in the treatment group.[25]

Therefore, we select comparison municipalities for the control group using the synthetic control estimator developed in Abadie, Diamond, and Hainmueller (2010) (hereafter ADH).[26] For each municipality with a college opening, we use the ADH method to construct an optimal synthetic control group, with pretreatment trends in the outcome variable matching those of the treated municipality. 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 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

25. Table 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.

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

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.[27] 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 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.[28]

3.2. *Effects of the Reform on Skill Compositions and Wages*

3.2.1. Main Results. Figure 1 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 A5 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 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 1, the year of the reform for each municipality is normalized to period zero. For every treatment or control municipality, we compute

27. In cases where the pre-reform period is less than five years, Z_{it} includes the pretreatment trends in the outcome variable in all available years prior to the reform.

28. We 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). Although we think the municipality level is a reasonable approximation of local labor markets and ensures comparability with previous empirical work that relies on geographical segmentation of the Norwegian labor market (see e.g. Black, Devereux, and Salvanes (2005); Akerman, Gaarder, and Mogstad (2015)), we check how spillover effects can affect our results. To this end, we exclude the municipalities within 30 km of the treated municipality from the “donor pool”, and re-estimate how college openings affect relative wages and supply using the synthetic control method. We find that the results do not materially change, demonstrating the estimated effects of the reform are not driven by comparison between the treated municipalities and nearby municipalities where there is most commuting (the full results are available upon request).

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 1 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 every 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 B.[29]

The top-left panel of this figure shows that, compared with the synthetic controls, labor markets with a new college experience an increase in the supply of skilled young 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.[30] In contrast, the reform has little impact on the skill composition among workers aged 35 years or more for the first 10 years following the reform. The 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.

The lower panel of Figure 1 shows estimates of the effects of the reform on the relative earnings of skilled workers, by age groups. Immediately after the college opening, the relative earnings of young skilled workers are slightly higher in the control than in the treatment group. Towards the end of the period we study, the relative earnings of skilled workers in the treatment group increase above those in the control group. Among older workers, the relative earnings of skilled individuals increase substantially following a college opening (this pattern is also seen for absolute earnings of skilled workers, as shown in Figure A5 in the Appendix).[31]

29. The treatment effects reported in the synthetic control figures such as Figure 1 are constructed by averaging municipality specific the treatment effects, one for each of the treated municipalities. The composition of the treated municipalities varies by pre-reform years, which may impact our results. For instance, in Figure 1, the relative wages and skill shares in the treated municipality and the synthetic control in period -5 are computed only for municipalities which had a new college opened from year 1972 and onwards, whereas in period -2, all municipalities that had a new college are included. As a robustness check, we also re-estimate the effects of the reform on wages and skill composition using the synthetic control method where we match the pre-reform outcomes using data only up to 2 years prior to the reform (for all treated) municipalities, where the composition of the treated municipalities is fixed in each of the pre-reform years (see Appendix Figure A6). We find effects on relative wages and skill shares that are very similar to what we report in Figure 1.

30. When 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 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.

31. Norway is characterized by a centralized wage bargaining system where wages are partially set at the national level through the central employer's and employee's organizations. This implies that there is a stronger degree of wage compression than in most other countries (Moene and Wallerstein, 1997; Kahn, 1998). Moreover, in the 1970s the Norwegian government more actively took part as a third party in the central negotiations with the aim of reducing the nominal increases in wage rises, resulting in very little (if

A reasonable interpretation of these findings is the following. The drop in the relative wages of young skilled workers immediately after the reform is a consequence of increased supply putting downward pressure on wages. However, the fact that, in spite of the rise in supply, the relative wages of skilled young workers in treatment areas eventually rises above that of control areas, suggests that the demand for these workers increased more rapidly in treatment than in control areas. This could be possibly caused by an endogenous increase in investment in skill biased technologies by firms in treatment areas, since the abundance of skilled workers may have made the use of these technologies more profitable.[32] This would be consistent with our findings that older skilled workers experience a stronger increase in their relative wages in treatment than in control areas. Older skilled workers are shielded from any downward pressure on their earning due to the 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. In Section 3.3 we develop in more detail this interpretation of our findings.

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 B). 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 municipality specific p-values for the test that the treatment effect is zero. However, it is simpler to present a single p-value for the treatment effects averaged across all treated municipalities. We begin by randomly drawing 50 placebos (with replacement) for each treated municipality, and average them across all treated municipalities. We then calculate p-values based on the distribution of the treatment effects from these aggregated placebos.

Figure 2 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 1), with the observed treatment assignment. The implied p-values for each of the actual gaps in each year, i.e., the

any) real wage growth (Aanensen, 2010). Importantly though, since we do not rely on the national variation as a source of identification, these national changes are in principle differenced out in our estimates.

32. One 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 D for details.

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 1 and 2 (Figure A7 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 C, 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 aggregate skill and relative wages across the two age groups are qualitatively similar to the synthetic control estimates.

3.2.2. STEM vs non-STEM Colleges. Out of the 15 new colleges in our study, 9 had majors in STEM fields (we call these STEM colleges). 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, and conversely, STEM graduates may provide the greatest incentive to firms for upgrading their technology. 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 3 and 4 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 1 are driven exclusively by those regions where a STEM college was established.[33]

33. When we aggregate the synthetic control estimates from each individual reform municipality, we use the number of plants in the municipality as weight. For municipalities with a non-STEM college, the average number of plants is smaller than the municipalities with a STEM college. Therefore, STEM

Labor markets where there was an opening of a non-STEM college did not experience substantial wage changes relative to control labor markets. The relative wages of skilled old workers move similarly in control and treatment areas. Similarly, in the years immediately following the opening of a college, the relative wages of skilled young workers do not follow distinct paths in treatment and control areas, perhaps because the inflow of skilled workers is not large enough to make a difference at the very beginning, or because it takes some time for wages to adjust. However it is remarkable that several years after the reform the relative wages of skilled young workers are much lower in treatment than in control areas, probably due to the downward pressure on wages resulting from an increase in the supply of skilled workers in the medium run (and the absence of a technology response). The results concerning earnings levels (as opposed to relative earnings), which show an increase in the earnings of skilled workers in areas where there was an opening of a STEM college, are presented in Appendix Figures A8 and A9.

3.3. Separating Supply and Demand Factors

3.3.1. The model setup. We use the model in Card and Lemieux (2001) to decompose the differences in trends in skill- and age-specific wages between reform and non-reform areas (reported in Figure A5) into supply and technology factors. Assume that aggregate output in period t and labor market D depends on two CES sub-aggregates of skilled (college) and unskilled (non-college) labor:

$$Y_{t(D)} = (\alpha_{t(D)}(a_{t(D)}S_{t(D)})^{\rho} + (1-\alpha_{t(D)})(b_{t(D)}U_{t(D)})^{\rho})^{\frac{1}{\rho}} \tag{1}$$

and

$$S_{t(D)} = [\sum_j \beta_j^s S_{jt(D)}^{\eta}]^{\frac{1}{\eta}}$$
$$U_{t(D)} = [\sum_j \beta_j^u U_{jt(D)}^{\eta}]^{\frac{1}{\eta}}$$

The gross elasticity of substitution between different age groups j with the same level of skill is $\sigma_A = 1/(1-\eta)$ where $\eta \in (-\infty, 1)$. Workers of different ages 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.[34] $\sigma_E = 1/(1-\rho)$, where $\rho \in (-\infty, 1)$, is the elasticity of substitution between skilled and unskilled workers and substitutes and complements are defined as above

colleges carry a larger weight when it comes to the overall effects, contributing to the similarities between Figure 3 and Figure 1.

34. We have assumed that the elasticity of substitution across age groups is the same for both skill groups. In Appendix Section E.1, we show how the model can be extended to allow for separate elasticities across skill groups. Estimation results from the extended model are very close to the model we present in this Section.

(these elasticities of substitution determine how changes in technology and changes in the supply of skill across cohorts affect the demand for skill and the relative wages of skilled workers). β_j^s and β_j^u are efficiency units of skilled and unskilled labor of age group j, respectively.

Note that this formulation of the CES production function allows factor-augmenting technologies to affect the productivity of workers through the efficiency units of labor.[35] $a_{t(D)}$ and $b_{t(D)}$ represent skilled and unskilled labor augmenting technological change, and $\alpha_{t(D)}$ can be interpreted as indexing the share of work activities allocated to skilled labor (Autor, Katz, and Kearney, 2008). Skill-biased technical change involve increases in $a_{t(D)}/b_{t(D)}$ or $\alpha_{t(D)}$.[36]

Let $D=1$ denote the treatment group and $D=0$ denote the synthetic control group. D can affect inputs ($S_{jt(D)}$ and $U_{jt(D)}$) and the technology parameters ($a_{t(D)}$, $b_{t(D)}$, $\alpha_{t(D)}$) in the post-reform periods ($t \geq 1$). The technology response may be an endogenous response to changes in $S_{jt(D)}$ and $U_{jt(D)}$. In period 0 ($t=0$; pre-reform period), the treated group and the synthetic control have identical labor inputs and productivity parameters. After period 0, labor inputs and technology parameters may differ between the treated group and the synthetic control.

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

$$\log \frac{w^s_{jt(D)}}{w^u_{jt(D)}} = \underbrace{[\log(\frac{\alpha_{t(D)}}{1-\alpha_{t(D)}}) + \rho \log(\frac{a_{t(D)}}{b_{t(D)}})]}_{\equiv \log(\frac{\theta_{st(D)}}{\theta_{ut(D)}})} + \log \frac{\beta_j^s}{\beta_j^u} - \frac{1}{\sigma_E} \log \frac{S_{t(D)}}{U_{t(D)}} - \frac{1}{\sigma_A} (\log \frac{S_{jt(D)}}{U_{jt(D)}} - \log \frac{S_{t(D)}}{U_{t(D)}}) \tag{2}$$

Note that the aggregate supply of skill ($S_{t(D)}$ and $U_{t(D)}$) is unobserved and depends on age specific skill supplies ($S_{jt(D)}$ and $U_{jt(D)}$) and the parameters in the sub-aggregate CES production function (η, β_j^u, β_j^s).

Let $d\log(S_{1t}/U_{1t}) \equiv \log(S_{1t(1)}/U_{1t(1)}) - \log(S_{1t(0)}/U_{1t(0)})$ denote an increase of the relative supply of young workers in the treatment group relative to the control group. Equation (2) shows that such an increase in the relative supply of young

35. One 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 4, 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.

36. In the literature, an increase in $a_{t(D)}/b_{t(D)}$ represents intensive SBTC and an increase in $\alpha_{t(D)}$ represents extensive SBTC (Johnson, 1997). Intensive SBTC makes skilled labor more productive at the tasks it already performs, without replacing the tasks of the unskilled labor. Extensive skill-biased technological change can be interpreted as changes in production processes such that skilled workers are profitably employed in some jobs that unskilled workers used to do. This type of technical change is a productivity gain by one factor and a loss by another.

workers can have two direct effects. The first comes from changes in the aggregate relative supply, $-(1/\sigma_E)d\log(S_t/U_t)$. This effect is the same for young and old workers and depends on the elasticity of substitution between skilled and unskilled workers. The second comes from the direct negative effect on relative wages among young workers, given by the term $-(1/\sigma_A)(\log(S_{jt(D)}/U_{jt(D)})-\log(S_{t(D)}/U_{t(D)}))$. For old workers, a rise in the relative supply of young workers may raise their relative wage (by $(1/\sigma_A)d\log(S_t/U_t)$). In the extreme case where young and old workers are perfect substitutes within skill groups (i.e., $\sigma_A \to \infty$), then the effect on young workers is identical to the effect on old workers.

The above two effects are standard, and due to supply factors. Our model allows for a third effect of the reform on relative wages, due to technology change. Our main goal is to estimate the sequence of $d\log(\theta_{st}/\theta_{ut}) = \log(\theta_{st(1)}/\theta_{ut(1)}) - \log(\theta_{st(0)}/\theta_{ut(0)})$, using the data generated in Figure 1.[37] These parameters identify the differential technical change in the treated group relative to the control. Given that we assume that the only direct effect of the reform on the economy is a change in the relative supply of skill, the only way technology could have responded in this model is if firms invested in technology upgrading as an endogenous response to changes in skill supplies. Therefore, we say there is evidence of *endogenous* skill-biased technical change if the treated group experienced more accelerated skill biased technical change than the synthetic control, for the relative supply of skills increased more in the treated group than the synthetic control.

It is however possible that the reform could have affected technical change for reasons unrelated to the supply of skill. 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. In Section 3.3.5, we discuss why this and other alternatives can be ruled out.

3.3.2. Identification and estimation strategies. When estimating equation (2), we face two empirical challenges. First, to credibly 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 $\theta_{st(D)}/\theta_{ut(D)}$. Using college openings alone as an exogenous shock is not enough for

37. We 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.

separately identifying these two mechanisms. Below explore additional assumptions which allow us to separate these two effects.[38]

We estimate equation (2) in two steps, using the data generated in Figure 1.[39] Below we discuss the identification arguments in each of the two steps, leaving details of the empirical implementation in Appendix Section E.

In the first step, σ_A (the gross elasticity of substitution between different age groups j within a given skill group) is identified from exogenous shifts in age-specific skill supplies. In particular, we explore exogenous changes in supply (and, therefore, identify the demand curve) by using differences in relative supplies within age (j) and across treatment groups (D) and correlate them to differences in relative wages.[40] Given the estimate of σ_A, the efficiency parameters β_j^s and β_j^u (which are assumed to be invariant to D) are estimated using the equations derived by equalizing the marginal product of labor with the wage for each combination of age and skill groups (as in Card and Lemieux (2001)).

In the second step of the estimation, we use data from both the treated group and the control group to identify the effects of college openings on technology change. It is here that we face the challenge of separately identifying σ_E and endogenous technical change parameters. Based on equation (2), and equipped with the estimated σ_A and aggregate supplies, we can rewrite the model as:

$$\log \frac{w^s_{jt(D)}}{w^u_{jt(D)}} + \frac{1}{\sigma_A}(\log \frac{S_{jt(D)}}{U_{jt(D)}} - \log \frac{S_{t(D)}}{U_{t(D)}}) = \delta_{t(D)} + b_j - \frac{1}{\sigma_E} \log \frac{S_{t(D)}}{U_{t(D)}} + e_{jtD} \quad (3)$$

where b_j are age-group dummies, and $\delta_{t(D)} \equiv \log(\theta_{st(D)}/\theta_{ut(D)})$ represents the relative technology efficiency which is specific to each year and treatment group. Our parameters of interest are the sequence of $\delta_{t(D)}$ and σ_E. Following much of the literature on this topic, the relative technology efficiency, we model $\delta_{t(D)}$ as a linear trend (although we could relax this assumption), interpreted as skill-biased technical change (Katz and Murphy, 1992):[41] $\delta_{t(D)} = \delta_0 t + \delta_1 (t \times D) + \delta_2 D$. The trend in

38. This problem is ignored in standard papers on wage inequality such as Card and Lemieux (2001) and several others in this lage literature because of the implicit assumption that technological change is exogenous.

39. We 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.

40. In the production function, technology is assumed to operate uniformly across age groups: β_j^s and β_j^u do not vary with t. With this assumption, σ_A can be identified in the data using within municipality and skill movements in age-specific supplies as we have shown. Note that our approach can still identify σ_A from the data even if we allow for age-specific technical change, as long as the age-specific technical change is exogenous (i.e., if they evolve over time in the same way between the treatment group and the control group).

41. The 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 $\theta_{st(0)}/\theta_{ut(0)} = \theta_{st(1)}/\theta_{ut(1)}$ up to the first M years after the college opening, we can identify σ_E and σ_A

technical change is allowed to vary with treatment. δ_0 represents skill-biased technical change in the synthetic control group, whereas δ_1 represents the incremental skill-biased technical change taking place in the treated group. A positive δ_1 implies endogenous technical change.

In this equation, $S_{t(D)}/U_{t(D)}$ could be correlated with e_{jtD}, leading to biased estimates of all the parameters. More importantly, any exogenous variation in skill supplies is no longer a valid exclusion restriction for $S_{t(D)}/U_{t(D)}$ in this equation because it also has a direct effect on technical change, $\delta_{t(D)}$, through the channel we emphasize in the paper.

To address this issue, we proceed with two different identification strategies. One is to use an external estimate of σ_E from the literature to back out the technical change parameters. An advantage of this approach is that we can experiment with a range of plausible values of σ_E to gauge the amount of technical change that is needed to match our data.

A second idea is to make one additional timing assumption in order to identify both σ_E and $\delta_{t(D)}$ from our data. A reasonable possibility is to assume that $\theta_{st(D)}/\theta_{ut(D)}$ does not evolve differentially with D in the years immediately following the reform. Formally, this means that $\theta_{st(0)}/\theta_{ut(0)} = \theta_{st(1)}/\theta_{ut(1)}$ 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 in equation (E.5) for fixed $\theta_{st(D)}/\theta_{ut(D)}$, by relating differences in relative wages to (exogenous) 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 $\theta_{st(D)}/\theta_{ut(D)}$, $t > M$.[42]

Our timing assumption can be motivated by existing models of technology adoption and innovation in the literature. In Appendix Section F, we describe a model of endogenous technology adoption following Acemoglu (2007) and explain how the model can imply that there is a threshold for the supply of skilled labor beyond which technical change takes place, and below which it does not (because when there is a fixed cost of technology adoption, and firms only have incentive to adopt the new

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 M years after the opening of the college).

42. Another intuitive idea would be to use the older workers in the years immediately after the reform to identify $\log(\theta_{st(1)}/\theta_{ut(1)}) - \log(\theta_{st(0)}/\theta_{ut(0)})$ because they did not experience increases in $S_{jt(D)}/U_{jt(D)}$ until much later. Given $\log(\theta_{st(1)}/\theta_{ut(1)}) - \log(\theta_{st(0)}/\theta_{ut(0)})$, 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 $S_{jt(D)}/U_{jt(D)}$ 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.

technology when return exceeds cost). This provides a theoretical justification for our empirical identification strategy.[43]

Empirical evidence on the timing of technology change after the reform provides further guidance to our timing assumption. This is more challenging because our R&D data do not allow us to conduct an event-study analysis to identify the timing of the adoption/innovation (see Section 5 for details). Nevertheless, investment in new equipment is likely to be complementary with high skilled jobs, which can be regarded as a proxy for technology adoption. Figure 6 shows the effects of the reform on investments in different types of new equipment: (i) machinery and equipment, (ii) machinery and equipment but excluding transport equipment (e.g. cars), and (iii) machinery, equipment and production facilities, where production facilities include infrastructure for production.[44] Appendix Table A4 shows the corresponding p-values from permutation tests. For all three measures, we find a positive effect on investments in new equipment in the treated municipalities at the time of the reform, which grows over time, and only becomes statistically different from zero from around 10 years after the reform. For the measure including production facilities, there is some weak evidence that the positive effects emerge earlier than for measures involving only machinery, which could be due to firms first upgrading production facilities before investing in new machines. The fact that these impacts are all relatively small (and statistically not different from zero) immediately after the reform provides some justification of our M-year lag identifying assumption, where substantial technology effects may only emerge after M years of the reform.

3.3.3. Can the model explain the movement of wages and supply without endogenous skill-biased technical change? Can the model explain the movement of relative wages and supply shifts shown in Figure 1, without any endogenous skill-biased technical change? Suppose that there is no endogenous SBTC, which means that $d(\theta_{st}/\theta_{ut}) = 0$, or equivalently, $\theta_{st(1)}/\theta_{ut(1)} = \theta_{st(0)}/\theta_{ut(0)}, \forall t$. Then, take the difference between treated group and the synthetic control within age groups, we get

$$d(\log \frac{w^s_{jt}}{w^u_{jt}}) = -\frac{1}{\sigma_E} d(\log \frac{S_t}{U_t}) - \frac{1}{\sigma_A} d(\log \frac{S_{jt}}{U_{jt}} - \log \frac{S_t}{U_t}), j = \{1,2\} \tag{4}$$

43. Other models, including Acemoglu (1998) and Beaudry and Green (2003) also show that an endogenous change in technology only takes place when the supply of skilled workers increases above a certain threshold. For instance, in Acemoglu (1998), new technologies are invented, and inventors devote more effort in the invention of skill-complementary technologies when there are more skilled workers (because when there are more skilled workers, the market for skill-complementary technologies is larger and the inventor is able to obtain higher profits). Given the fixed cost of technology invention, technology innovation will only take place when there is sufficient number of high skilled workers in the market. Therefore, there may be periods of increasing supply of skill without a corresponding change in technology.

44. Note that these variables measure the investments in new equipment—costs of repairing existing equipment are excluded. Measure (ii) is the narrowest measure for machinery and perhaps the closest to equipment used directly in the production process that we are able to obtain from our plant-level data (we do not have data on ICT equipment in the period under study).

For old workers ($j = 2$), let us assume that $d(\log(S_{2t}/U_{2t})) = 0$, which is consistent with our empirical results for the first 10 years following the reform. In these circumstances, in order to get positive wage effect on older workers as shown in Figure 1, we need that $\sigma_E > \sigma_A$, which is possible, but inconsistent with estimates of σ_E and σ_A in the literature (e.g., Card and Lemieux (2001)).

Our data also suggests that, for young workers, $d(\log(S_{1t}/U_{1t}) - \log(S_t/U_t))$ and $d(\log(S_t/U_t))$ are both positive.[45] Therefore, without any endogenous technical change, the model predicts a decrease in the relative wage of young workers, because both terms on the RHS of equation (4) are negative for young workers. This may be consistent with the data immediately after the reform, but not with the evidence for the subsequent years.

3.3.4. Estimation results. We start by presenting results where we fix σ_E and estimate the implied rate of technical change in treatment and control areas. Figure 5 shows the impact of a college opening on the resulting rate of endogenous technical change (the incremental trend in the unexplained relative wages in the treated group), for different values for the elasticity of substitution between skill groups. The less substitutable skilled and unskilled labor is, the higher is the rate of endogenous technical change implied by the model. Notice also that even if skilled-unskilled labor are highly substitutable, the implied endogenous technical change is still estimated to be positive and significantly different from zero. In Card and Lemieux (2001), the estimated σ_E is between 2 and 2.5. At these values, the implied impact of endogenous technical change on the relative wage of skilled workers is slightly over 0.005 (0.5%) per year.

Table 2 presents the estimated parameters resulting from models where we use the second identification strategy, relying on a timing of response assumption. 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 the treated group beginning two years after the opening of a new college. We estimate the interaction between the linear trend and the treatment dummy to be positive and statistically significant, showing that the reform leads to an incremental increase in the relative demand for high-skilled workers. Holding relative supplies fixed, a college opening in the municipality leads to an increase in the growth rate of the relative skilled wage of 0.8% per year. The implied elasticity of substitution between college and non-college labor is slightly below 2.[46]

In subsequent columns, we show that our results are robust to alternative assumptions about the value of M. In columns (2)-(4), we set $M = 3, 4, 5$, respectively. Each of these alternatives gives very similar conclusions to our baseline

45. S_t/U_t is constructed given the estimated σ_A, $\log(\beta_j^s)$ and $\log(\beta_j^u)$, as explained before. Note that if young and old workers are perfect substitutes, then $d(\log(S_{1t}/U_{1t}) - \log(S_t/U_t)) = 0$.

46. By 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).

specification: 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 A10 that the estimates of the baseline model of equations (E.1) and (E.5) provide a reasonably good fit to the skill premiums of the treated group and synthetic control group.

In column (1) of Appendix Table A5, we report the estimated parameters from the first step of the estimation, where we estimate σ_A using exogenous variation in relative supply. The year dummies show a pattern of steeply rising relative returns. The estimates imply an elasticity of substitution between young and old workers of about 3, illustrating the importance of considering imperfect substitutability between young and old workers for a given skill group.[47]

Our results are consistent with the model of endogenous technology adoption in Acemoglu (2007).[48] In Appendix Section F, we describe this model in its simplest case. In that model, firms in each market have access to the same set of factor-augmenting technologies and choose the type of technology it wants to adopt together with skilled and unskilled labor inputs. When skill levels are low and the relative cost of adopting the skilled-biased technology is large enough, initial optimal choice of technology is the least skill-biased one. When the supply of skilled workers in the market increases, skilled wages initially decrease holding technology choice constant. As the supply of skilled workers keeps increasing, eventually the marginal benefit from technology adoption exceeds the marginal cost and firms switch the choice of technology to the skill-biased one. 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. As discussed in Acemoglu (2007), this sort of dynamics may lead to a positive relationship between the quantity of skilled input and the wage of skilled workers, or a long-run, upward-sloping demand for skill.[49]

47. By 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 however that the age groups used in their paper are more disaggregated (five-year age bins defined from ages 26–60 years) than ours. It is also plausible to expect low elasticity of substitution between age groups, in a labor market where there is little competition between workers of different age groups. As described in Huttunen, Møen, and Salvanes (2011), the seniority rule is an important feature of the Norwegian labor market. This could be one possible explanation for the fact that the elasticity of substitution between age groups (σ_A) that we estimated is at the lower end in the literature (as reported in Appendix Table A8).

48. 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.

49. Another implication of our results is that standard estimates of supply effects and skill biased technical change in the literature may need to be amended. In Appendix Section G we have estimated a standard Card and Lemieux (2001) type model using national data for Norway. We show that incorporating our estimate of endogenous technical change in the empirical model halves the estimate of σ_E, and more than doubles that rate of skill biased technical change between 1967 and 1990.

3.3.5. Ruling out alternative explanations. Could there be other reasons, not related to endogenous technical change, that potentially explain the evolution of relative wages and skill supplies as observed in our data? One potential explanation is that, in addition to shifting the local skill compositions, these colleges had a direct impact on innovation. However, in the period that we study, these colleges did not produce patents or engage in R&D activities that may directly affect technological change in the private sector. Appendix Figure A4 shows that the fraction of higher education R&D expenditures taking place in regional colleges in the 1970s was extremely low (1%-3%).[50] In addition, although there is no reliable patent data dating back to the 1970s, patent data from the 1990s shows that universities are far more important for generating new patents than regional colleges (see Appendix Figure A12), and this is likely to be even more exacerbated in the 1970s. This is in contrast to new universities in other countries which have been found to create strong innovation spillovers to the private sector.[51] Therefore, in our context, the role of these new colleges was only to provide college education and not research.

Even though these new colleges did not engage in research activities, one might still be concerned that the opening of a college may increase the relative demand for high-skilled workers since many positions at the new colleges might require university educated workers, especially those with prior experience in the sector. If this is the case, the widening of the skilled wage gap over time in reform municipalities could be explained by the growth of the new colleges, and its resulting impact on the demand for skilled workers. We investigated this hypothesis by documenting, using decennial census data, the relative employment size of the college sector in the municipalities with a new college. In the 1980 census, only 1.93% of total employment in reform municipalities is in the “Universities and colleges" sector. Among high-skilled workers, only 2.95% worked as college lecturers in these colleges.[52] Therefore, any direct demand effect from college openings is likely to be very small. It is also unlikely that the construction of the regional colleges caused a significant increase in the demand for other skilled services, given how small they were relative to the overall size of the economy.[53]

We also examined the inflow of skilled workers into the municipality when a new college is constructed using an event-study. If these new colleges begin to hire many skilled workers, then we would expect to see more skilled workers moving into the

50. In the 1990s (outside our period of study) some of these colleges were converted into universities, but the expenditure share in R&D remains small.

51. See, among others, Kantor and Whalley (2014), Toivanen and Väänänen (2016), and Andrews (forthcoming).

52. Among all low-skilled workers residing in the reform city, 0.03% worked as college lecturers in 1980. College lecturers are defined using the Norwegian Standard Classification of Occupations 1965 (061–University professors and readers and 062–Other university lecturers).

53. Relative small colleges are able to produce large increases in the stock of skilled workers because they produce a constant flow of graduates over several years, which increasingly contribute to the existing stock of college educated workers.

reform municipality from elsewhere, especially in the first few years when the college expands. Appendix Figure A13 shows that there is no differential change between the treated and control municipalities in terms of inflow of skilled workers (as a fraction of the population size). Overall, we believe that the employment size of these colleges is too small to shift the local labor demand.

An increase in the supply of young skilled workers could generate an increase in wages of old skilled workers if they were somewhat complements in production. This was already discussed above, when we say that, taking for example equation (4), for this results we would need that $\sigma_E > \sigma_A$ (stronger substitutability between workers of different skill groups than between workers of different ages), which is at odds with the estimates in the literature showing precisely the opposite. Appendix Table A8 reports estimates from some of the central papers in the literature with estimates of σ_E and σ_A. If we take the estimates of Card and Lemieux (2001) or several other papers, keeping technology fixed, an increase in the supply of skilled young workers leads to a fall in the wages of skilled old workers. Furthermore, as explained above, it is impossible to use such a model (and the absence of endogenous technical change) to also justify an eventual increase in the wages of skilled young workers.

The model that we have presented so far has a single sector in a particular local labor market. A rival model is a trade model, where there are multiple sectors in the local market together with workers who are mobile across different sectors. With many sectors, each with a different skill intensity, an increase in the supply of skilled workers can potentially lead to a rise in output in the skilled-intensive sector, and more generally, a change in the local output mix (across sectors with different skill intensities), with no obvious implications for wages.[54]

In Appendix Figure A14, we compare the shares of total employment, output and plants in skill-intensive industries between the treatment and the control group. Skill-intensive industries consist of industries with above-median shares of skilled workers in 1967. We do not find large or significant differences between the treated and control group in terms of output/employment/plant compositions, suggesting that the skilled sector did not expand more than the unskilled sector immediately after the reform.[55] Therefore, it is unlikely that changes in sectoral composition are substantially affecting our results. Furthermore, in our production function analysis below, we always control for industry fixed effects. This implies that we uncover technical change responses in the firm-level analysis which are identified only from within-industry variation (as opposed to between-industry comparisons).

54. In general, the trade model implies that the *economy-wide* relative labor demand curve is horizontal because relative wage is determined only by output prices and TFP (the Rybczynski Theorem).

55. The p-values of the differences are reported in Appendix Table A9

4. Firm-level Evidence: Structural Estimates of the Production Function

In Section 3, we used labor market data to show that the relative wages 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 the 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, using firm level data on inputs and outputs.

Therefore, we do not need to rely on the assumption of competitive labor markets to learn about the impacts of the reform on the productivity of skilled and unskilled labor. This is potentially important because centralized wage bargaining was important in Norway during the period we study, which means that fluctuations in productivity may not fully pass onto wages (at least in the short run). The effects of the reform on labor productivity could well 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 occurs in spite of an increase in the amount of skilled labor used by firms. Again, as discussed in Section 3, our conjecture is that these results are driven by endogenous technical change responding to an increase in the abundance of skilled labor.[56]

4.1. Empirical Strategy

As in Section 3.3, we assume that the production function depends on skilled and unskilled labor. In this section we will also allow it to depend on capital. Specifically, the production function for industry j in municipality c and year t takes the following

56. In Appendix Figures A15 and A16, 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.

form:[57]

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

where Y_{jct} is the total real value-added output at factor prices (before taxes). The three types of inputs used in production are: K_{jct}, the total real value of capital stock, S_{jct}, the total employment of skilled workers, and U_{jct}, the total employment of unskilled workers. Inputs, especially S_{jct} and U_{jct}, may depend on $\boldsymbol{d}_{ct}$, a vector of current and lagged reform indicators, defined by $\boldsymbol{d}_{ct} = \{d_{ct}^{\tau}\}_{\tau=0}^{R}$, where d_{ct}^{τ} 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 $\boldsymbol{d}_{ct}$. μ_{jct} is a productivity shock.

When discussing identification of the parameters of this model, and ignoring the distinction between workers of different ages (we come back to this point below), it is useful to rewrite skill-specific 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 relatively 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 to π_{jct} (skill composition of employment).

Our main objective is to study how the marginal product of skilled and unskilled labor depends on θ, which, in turn, depends on $\boldsymbol{d}_{ct}$. We assume that $\boldsymbol{d}_{ct}$ 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). As before, we need to face the problem that π_{jct} has a direct effect on Y_{jct}, but it also has an indirect effect (which is the focus of this paper) through θ_{ct}. To separate these two channels, we use one of the same assumptions discussed above: that θ_{ct} does not vary with π_{jct} in the first M years immediately following the opening of a college.

57. Note that the 3-digit industry classification we use is quite detailed, so the cells are narrow. Examples of 3-digit industries include: manufacturing of beverages (SIC 313); textiles (SIC 321); electrical apparatus and supplies (SIC 383); transport equipment (SIC 384). Altogether there are 22 industries within manufacturing, and in many cases (more than 10% of the cells) there is only one plant per industry-municipality-year cell. The average number of plants within a municipality-industry-year cell is 6.1, and the median is 3. In our production function analysis, we always use weights by the number of plants within industry-municipality-year cell. Therefore, our variations we exploit are closest to estimating production function from plant-level data.

We use a control-function approach to account for the endogeneity of skill shares, using the following econometric model:

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

where $\boldsymbol{d}_{ct}^{M} = \{d_{ct}^{\tau}\}_{\tau=M}^{R}$. The excluded instruments for π_{jct} are $\boldsymbol{d}_{ct} \setminus \boldsymbol{d}_{ct}^{M} = \{d_{ct}^{\tau}\}_{\tau=0}^{M-1}$. We approximate the control function with the following form (analogous to a simple series expansion):

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

If K_{jct} and L_{jct} could be assumed to be exogenous, we could implement the following estimator:

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

where $\widehat{v}_{jct}$ is the estimated residual from the first-stage equation for π_{jct}, which controls for the endogeneity of skill share in the production function. The residual ω_{jct} has zero conditional mean once $\widehat{v}_{jct}$ 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). Levinsohn and Petrin (2003) uses a structural model of an optimizing firm to derive the conditions under which the intermediate input demand function depends on the firm-specific state variables, including productivity shocks and capital. Under the assumption that the intermediate input demand is monotonic in the productivity shock for all capital, the intermediate input demand function can be inverted to yield a second control function for the unobserved productivity shock.[58] We refer to Appendix Section I for details of this control function approach. Our results are robust to the treatment of endogeneity of these other inputs.[59]

58. The control function approach of using intermediate inputs (the LP approach) to back-out productivity shock assumes that such productivity shocks are factor neutral (TFP shocks) and invertible. If we have a factor-biased productivity shock (a shock on the coefficients), using intermediate inputs alone may not be sufficient (Doraszelski and Jaumandreu, 2018). This limitation with the LP approach makes the experimental variations from the college opening reform (including the timing assumption) more valuable in terms of addressing the endogeneity of relative labor inputs (which respond to factor-biased productivity shocks).

59. Notice that we estimate production functions in the original dataset, as opposed to data generated by a synthetic control estimator, as in Section 3.3. 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 3.3, 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).

One practical difficulty is that, as discussed in Section 2.2,, we have annual measures of skill shares at the level of the municipality, π_{ct}, but not at the level of industry (or firm). We only observe this municipality-industry specific labour inputs from the decennial Census data. Industry level information on labour inputs is essential because the firm level data used in this section to estimate production functions is only available for manufacturing (and therefore we cannot estimate municipality level production functions). As a result we will be less ambitious in the specification of equation (5) than we were for equation (1), by ignoring the age distinction in the definition of labour aggregates. This distinction was illuminating when we focused on wages, and we could explore it because we had detailed data on labour inputs and wages of different groups. The estimation of even the simpler specification of the production function is then made possible by our use of an interpolation procedure to construct annual municipality level series for skilled and unskilled labour inputs in the manufacturing sector, which we describe in Appendix Section H.

In the above discussion on identification, we abstracted from any covariates. However, we need 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. Including them in a typical CES specification such as that of equation (1), augmented with capital and other covariates, is not practical. One alternative, following Kmenta (1967), is to linearize the CES aggregate (equation (1)) around $\rho = 0$ using a second-order Taylor expansion. In Section 4.2 and Appendix Section I, we consider other models and find our main results robust to alternative parameterizations, including a more flexible translog production function and a simpler Cobb-Douglas specification.

One potential concern is that there may be some other unobservable (e.g. at the industry level) that affects technology adoption and that is potentially correlated with skill intensity. If that's the case, then the interaction between the reform and labor inputs does not identify the impact of the reform on the marginal product of labour. In Section 4.2, we show that our results are robust to allowing for industry fixed effects in the productivity of skilled/unskilled labour, capturing permanent unobserved heterogeneity in the labor productivity across different industries (but the coefficient measuring the impact of the reform on the productivity of skilled labour does not depend on industry). Our results are robust to additionally allow pre-reform characteristics at the municipality-industry level to affect the labor productivity. Our identifying assumption is that, conditional on the covariates, there is no correlated unobservable in the parameters of the production function.

4.2. Estimation Results

In Table 3, we report the mean predicted differences in output elasticities of skilled and unskilled labour between treated and control group. Estimates come from the CES, the translog, and the Cobb-Douglas specifications of the production function, for different years following the reform. There are two columns for each type of production function. In the first one of each pair (columns 1, 3 and 5) we only account for the endogeneity of skill shares with the control function in equation (6), while in the second (columns 2, 4 and 6) we also account for the endogeneity of total employment with the control function in equation (I.7) in Appendix.[60]

The table has two panels. The top panel shows the impact of the reform on output elasticities of skilled labor 2, 5, 10 and 15 years after the reform, implied by each of the production functions, as well as the impact of the reform on the average annual growth in the output elasticity of skilled labor. The bottom panel provides estimates of the impact of the reform on the levels and growth of the output elasticity of unskilled labor. If we take, for example, the specification in column 1, five and ten years after the opening of the college, the output elasticity of skilled workers is 4.2 and 9 percentage points higher than it would have been in the absence of the reform, respectively. The corresponding numbers for unskilled workers are -0.033 and -0.092, five and ten years after the reform, respectively.

We find that the reform leads to an increase in the productivity of skilled labor and a decrease in the productivity of unskilled labor. Moreover, these impacts grow over time, presumably because the stock of skilled workers in the labor market is also growing over time, and firms are adjusting their technology accordingly. As mentioned above, the Cobb-Douglas specification is the one most restricting the substitution patterns across different inputs. However, allowing for more flexible substitution patterns in the remaining columns of the table does not change the magnitude of our estimates.

Overall, our results are consistent with both absolute skill-biased technological change (the productivity of skilled workers increases) and relative skill-biased technological 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.

Although we do not directly observe technology adoption from the firm-level data used in this section, Norwegian technology historians indicate that there was increasing use of ICT technology across different industries and improved organizational structure in the period we study, both of which may favor skilled workers at a cost of unskilled workers (in Section 5 we show impacts of the reform on R&D activities, from a different dataset). For instance, as in many other countries, automation (using computers) was introduced in the late 1960s and these technologies started to spread to many industries also in Norway starting in the 1970s and even more so in the 1980s (see Sogner (2002), and Wicken (1994) for overviews for Norway in this period, and for US evidence for instance Jovanovic and Rousseau (2005)). ICT technologies began being developed and spreading to different types of use in

60. Results from the first-stage control function equation are reported in Appendix Table A11.

the 70s and 80s (for an overview see Bresnahan (2010)).[61] The early ICT industries become important in the modernization of many sectors such as metal industries, ship-building, fisheries including land based processing of fish telecommunication and in the late 1970s and especially 1980s the oil sector.[62] For many industries the introduction of ICT technology was combined with the use of new form of organization of production and new types of management, leading to new forms of more flexible production.

Specification Checks

To separate supply and technology effects, we assume that there is no endogenous technological response immediately after the reform. To understand how robust our results are to our exclusion restriction, we estimate the Cobb-Douglas production function using different M-years as exclusion restrictions. Appendix Table A12 shows that the production function estimates are not affected by different assumptions of M. If any, using the first 5 years as exclusion restrictions implies slightly larger skill-biased technical change. Although using a larger value of M leads a stronger first stage because it leads to larger and more significant changes in supply, we find it reassuring that our main parameter of interest is not sensitive to the specific timing we use as exclusion restriction.[63]

Under an alternative set of assumptions (arguably less credible), we can allow the technology response to occur immediately after the reform. We present estimates of such a model in Appendix Table A13. 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 the productivity of both types of labor are unaffected by the reform in the short run. The estimated coefficients on the short-run reform

61. One of the largest producers of mini computers was located in Norway called Norwegian Data (Sogner (2002)). Combined with a group of engineers at the well established Norwegian Technical University (NTH), they develop computerized automatization systems for many process industries.

62. For instance, one of the industrial processes where computerized automation took place was in the metal producing sector, where the process of producing metals, for instance aluminum, was controlled by computerized systems (Sogner, 2002). Other sectors where these type of systems were introduced were the large hydro power industry producing electricity, the production of weapon which also was a quite large industry exporting for the international market, the ship-building industries, and shipping equipment industries. An example form the shipping industry is the systems developed to control the engines from the bridge of the ship as well as radar systems.

63. In order to assess the magnitude and the significance of the effect in the first M years on skill shares (our exclusion restriction), in columns (2)-(5) of Appendix Table A11, we report regressions where we distinguish the effects in the first M years vs. years beyond the first M years in the first stage regression. The coefficients on $D_{ct} \times (1 - P_{t(D)})$ identify the effects of the reform on skill shares in the first M years, where $M = 2, 3, 4, 5$ (corresponding to columns 2 to 5). As we include additional years as the exclusion restriction, the effects on skill shares become stronger. For instance, the first 5 years combined leads to strongly significant effect on skill shares, with a joint F-test of 4.5.

indicators interacted with skilled and unskilled labor are, both individually and jointly, insignificantly different from zero. Therefore, this 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 of the average annual growth in the output elasticities of skilled and unskilled labor for five other alternative specifications (see the Appendix Section I for details). We focus on the simpler Cobb-Douglas specification, although our results are similar regardless of which of three specifications of the production function we use. Columns (1) through (4) 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 and column (3) includes initial municipality-industry specific characteristics (from the 1960 census) in $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$. As discussed previously, results from these two columns are informative as to whether the interaction between the reform and labor inputs capture the impact of the reform on the marginal product of labour or be contaminated from some other unobservable (e.g. at the industry level) that is potentially correlated with skill intensity. Column (4) adds treatment group indicators, including a treatment group indicator and another indicator based on the timing of the treatment. Column (5) allows for the reform to impact the productivity of capital. The coefficients on labor inputs interacted with reform indicators are robust to the above specifications.

4.3. Quantifying the Technology Effects on Wages

To better visualize the quantitative implications of our estimates, in this section, we simulate the predicted marginal products of skilled and unskilled labor in the post-reform years. We compute the ratio of the two (or the difference in logs), and decompose changes in this variable into effects of changes in endogenous technology and effects of changes in supply.

Denote the log of the relative marginal product of labor in treatment group D by $\Delta_{jct(D)} \equiv \log MP^{s}_{jct(D)} - \log MP^{u}_{jct(D)}$, where $\log MP^{s}_{jct}$ and $\log MP^{u}_{jct}$ are the log of the predicted marginal product of skilled and unskilled labor, respectively (after taking into account the effect of the reform in each year t). In treated municipalities, the predicted change in $\Delta_{jct(D)}$ that is due to the reform in year t (after the year of reform) can be decomposed into technology effects and supply effects. Using the Cobb-Douglas specification in equation (I.2), we develop the following decomposition:[64]

$$\Delta_{jct(1)} - \Delta_{jct(0)} = \underbrace{(\log\beta_{1,ct(1)} - \log\beta_{2,ct(1)}) - (\log\beta_{1,ct(0)} - \log\beta_{2,ct(0)})}_{\text{technology effects}}$$
$$- \underbrace{\left[(\log S_{jct(1)} - \log U_{jct(1)}) - (\log S_{jct(0)} - \log U_{jct(0)})\right]}_{\text{supply effects}}$$

The advantage of using the Cobb-Douglas is that this decomposition can be written as this simple combination of coefficients and inputs.

Figure 7 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, M periods after the reform the technology effect dominates the supply effect, resulting in an increase in the level and growth of the relative marginal product of skilled labor (while immediately after the reform, before technological upgrading takes place in firms, the marginal product of skilled labor declines).

Notice that the predicted relative productivity 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 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).[65] 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.

64. $\log S_{jct(1)} - \log S_{jct(0)}$ 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}.

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

5. Additional Evidence from R&D Activities

Our results so far show that the opening of a college leads to an increase in the relative demand for skill, and an increase in the productivity of skilled workers. We interpret these findings 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.

Using firm-level R&D data, we investigate whether the college expansion reform led to an increase in R&D activities. The unit of observation in the R&D data is called "bransjeenhet," which consists of all plants of a firm with 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, we measure 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 dividing the total R&D activities by the number of plants with the same economic activity code.[66]

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:

$$\log(Y_{ct}) = \theta_c + \gamma_{s(c)t} + a_1 D_{ct} + \varepsilon_{ct} \tag{7}$$

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 variable taking value one if the municipality c has a college in year t; θ_c are fixed municipality effects; and $\gamma_{s(c)t}$ are county–year fixed effects. We consider two measures of R&D activities: i) firm expenditures in R&D activities and ii) man-years devoted to R&D activities.[67]

Table 5 reports our estimates of the parameters of equation (7). 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 a college opening increases total R&D expenditures by over 80%. This 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

66. The 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% of the firm-activity units with one or two plants and 52% of the firm-activity units located in a single municipality.

67. The 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, and, therefore, resulting in a smaller sample.

large positive impact on the man-years performing R&D activities in firms. Including municipality-specific linear trend makes the estimates less precise (due to the small sample used in this regression), but overall our estimates point to a large positive increase in the man-years performing R&D activities (which more than doubled) in treatment areas (relative to control areas) following the reform.

6. 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 large literature asks why SBTC was so pronounced during this period, and one possibility is 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 can cause 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 the 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 a new college. We use spatial and temporal variation in the availability of new colleges across local labor markets as a natural experiment to identify the impact of changes in the local supply of skill on local labor market outcomes. 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, education, employment and earnings.

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.

AANENSEN, T. (2010): "Lønnsutvikling for ansatte i skoleverket 1959-2008," .

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," *The Quarterly Journal of Economics*, 113(4), 1055.

ACEMOGLU, D. (2007): "Equilibrium bias of technology," *Econometrica*, 75(5), 1371–1409.

ACEMOGLU, D., AND J. ANGRIST (2001): "How large are human-capital externalities? Evidence from compulsory-schooling laws," in *NBER Macroeconomics Annual 2000, Volume 15*, pp. 9–74. MIT Press.

ACEMOGLU, D., AND M. DELL (2010): "Productivity Differences between and within Countries," *American Economic Journal: Macroeconomics*, 2(1), 169–88.

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.

ANDREWS, M. (forthcoming): "How Do Institutions of Higher Education Affect Local Invention? Evidence from the Establishment of US Colleges," *American Economic Journal: Economic Policy*.

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 (2017): "Life-cycle earnings, education premiums, and internal rates of return," *Journal of Labor Economics*, 35(4), 993–1030.

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(1), 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*.

BRESNAHAN, T. (2010): "General purpose technologies," in *Handbook of the Economics of Innovation*, vol. 2, pp. 761–791. Elsevier.

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.

CICCONE, A., AND G. PERI (2006): "Identifying human-capital externalities: Theory with applications," *The Review of Economic Studies*, 73(2), 381–412.

DAVIS, D. R., AND J. I. DINGEL (2019): "A spatial knowledge economy," *American Economic Review*, 109(1), 153–70.

DORASZELSKI, U., AND J. JAUMANDREU (2018): "Measuring the bias of technological change," *Journal of Political Economy*, 126(3), 1027–1084.

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.

HUNT, J., AND M. GAUTHIER-LOISELLE (2010): "How much does immigration boost innovation?," *American Economic Journal: Macroeconomics*, 2(2), 31–56.

HUTTUNEN, K., J. MØEN, AND K. G. SALVANES (2011): "How destructive is creative destruction? Effects of job loss on job mobility, withdrawal and income,"

Journal of the European Economic Association, 9(5), 840–870.
JOHNSEN, B. W. (1999): *Fra universitetsvisjon til høyskoleintegrasjon*. Kristiansand: Norwegian Academic Press.
JOHNSON, G. E. (1997): "Changes in earnings inequality: the role of demand shifts," *Journal of economic perspectives*, 11(2), 41–54.
JOVANOVIC, B., AND P. L. ROUSSEAU (2005): "General purpose technologies," in *Handbook of economic growth*, vol. 1, pp. 1181–1224. Elsevier.
KAHN, L. M. (1998): "Collective bargaining and the interindustry wage structure: international evidence," *Economica*, 65(260), 507–534.
KANTOR, S., AND A. WHALLEY (2014): "Knowledge spillovers from research universities: evidence from endowment value shocks," *Review of Economics and Statistics*, 96(1), 171–188.
——— (2019): "Research proximity and productivity: long-term evidence from agriculture," *Journal of Political Economy*, 127(2), 819–854.
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.
KMENTA, J. (1967): "On Estimation of the CES Production Function," *International Economic Review*, 8(2), 180–189.
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.
MOENE, K. O., AND M. WALLERSTEIN (1997): "Pay inequality," *Journal of Labor Economics*, 15(3), 403–430.
MORETTI, E. (2004a): "Estimating the social return to higher education: evidence from longitudinal and repeated cross-sectional data," *Journal of econometrics*, 121(1), 175–212.

MORETTI, E. (2004b): “Workers’ education, spillovers, and productivity: evidence from plant-level production functions,” *American Economic Review*, pp. 656–690.

MOSER, P., A. VOENA, AND F. WALDINGER (2014): “German Jewish émigrés and US invention,” *American Economic Review*, 104(10), 3222–55.

OTTOSEN-COMMITTEE (1966-1970): *Innstilling om videreutdanning for artianere m.v. fra komiteen til å utrede sprøsmål 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.

RAUCH, J. E. (1993): “Productivity gains from geographic concentration of human capital: evidence from the cities,” *Journal of urban economics*, 34(3), 380–400.

SOGNER, K. (2002): *En liten brikke i et stort spill: den norske IT-industrien fra krise til vekst 1975-2000*. Fagbokforl.

TOIVANEN, O., AND L. VÄÄNÄNEN (2016): “Education and invention,” *Review of Economics and Statistics*, 98(2), 382–396.

WICKEN, O. (1994): *Elektronikkentrepenørene*. Oslo: Ad Noram Forlag,(in Norwegian).

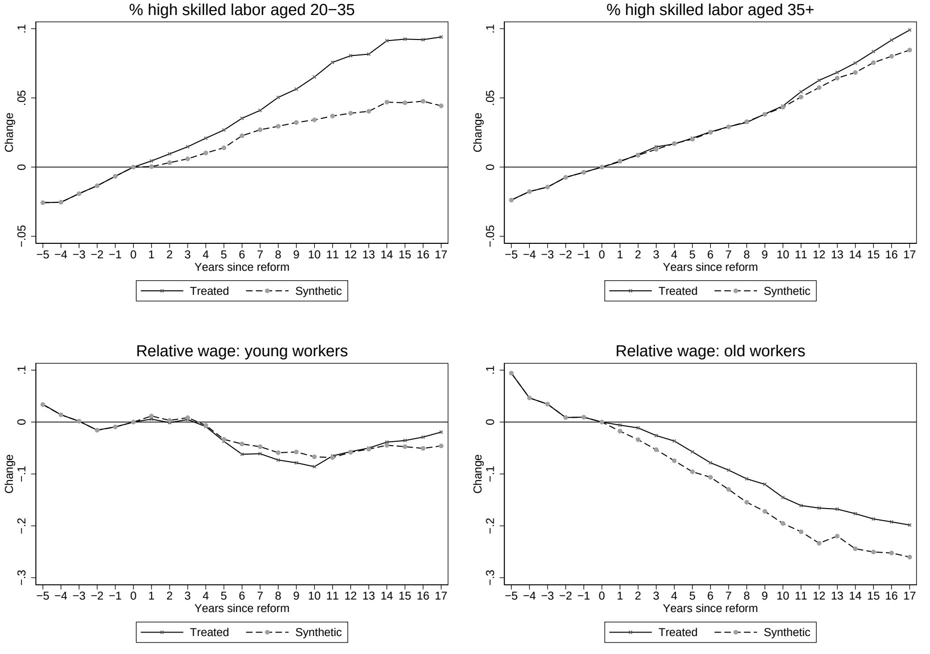

FIGURE 1. The Effects of the Reform on Relative Wages and Skill Compositions. This figure presents the synthetic control estimates on skill composition and relative wages of the workforce. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome 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.

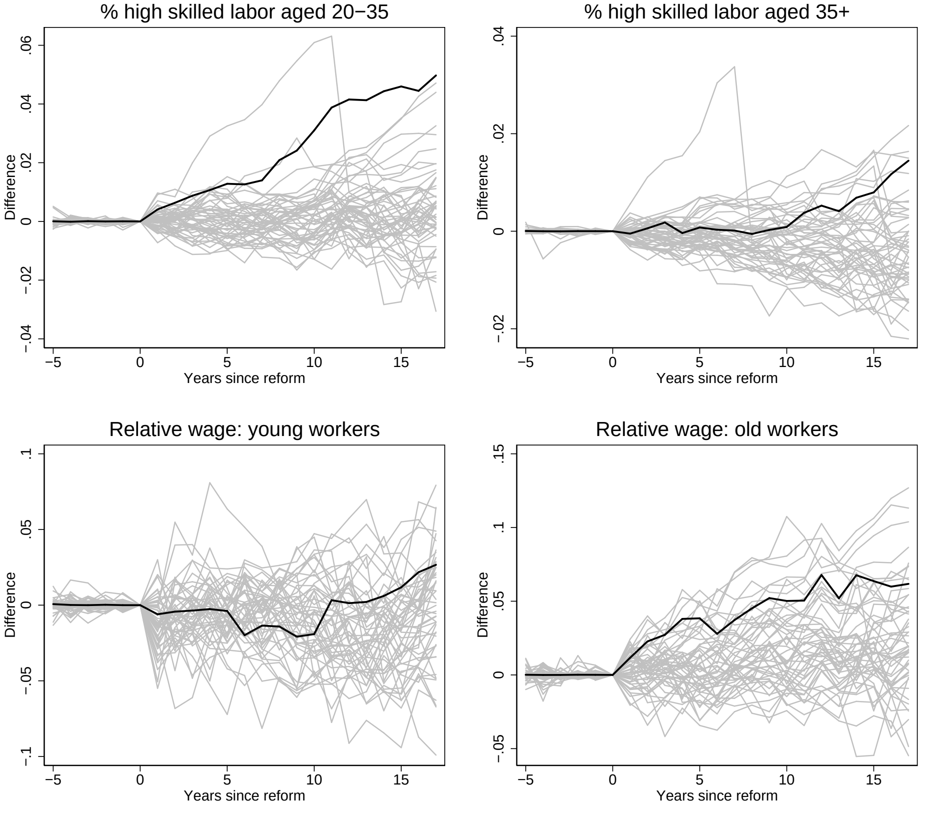

FIGURE 2. The Effects of the Reform on Relative Wages and Skill Compositions: Placebo Tests. This figure shows the results for the placebo test. The gray lines represent the estimated treatment effect from each placebo of the permutation test. The thick solid line denotes the treatment effect estimated 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.

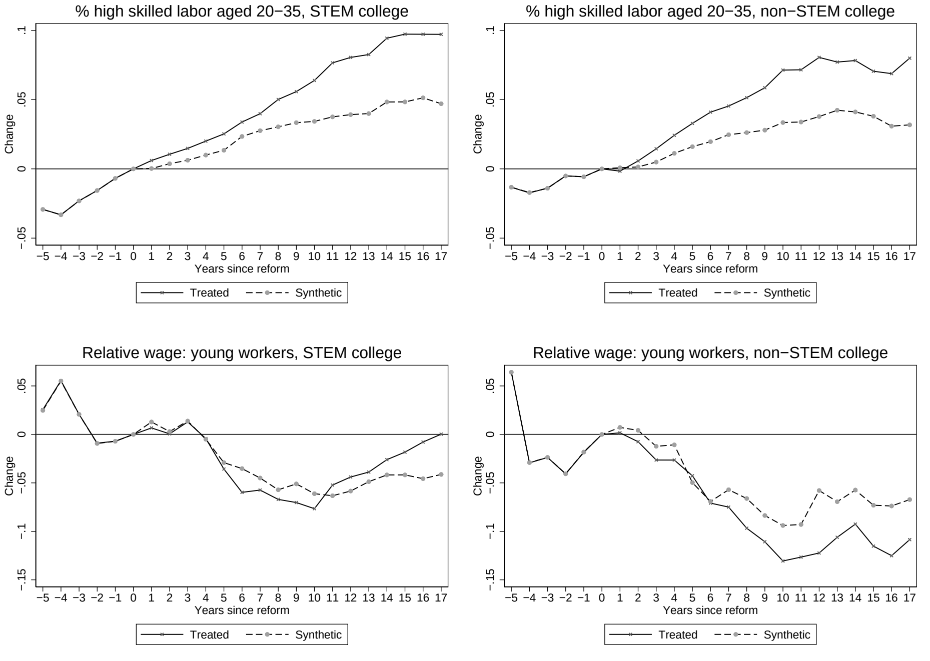

FIGURE 3. The Effects of the Reform on Young Workers: STEM vs Non-STEM Colleges. This figure presents the synthetic control estimates on skill composition and wages of young workers, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. 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.

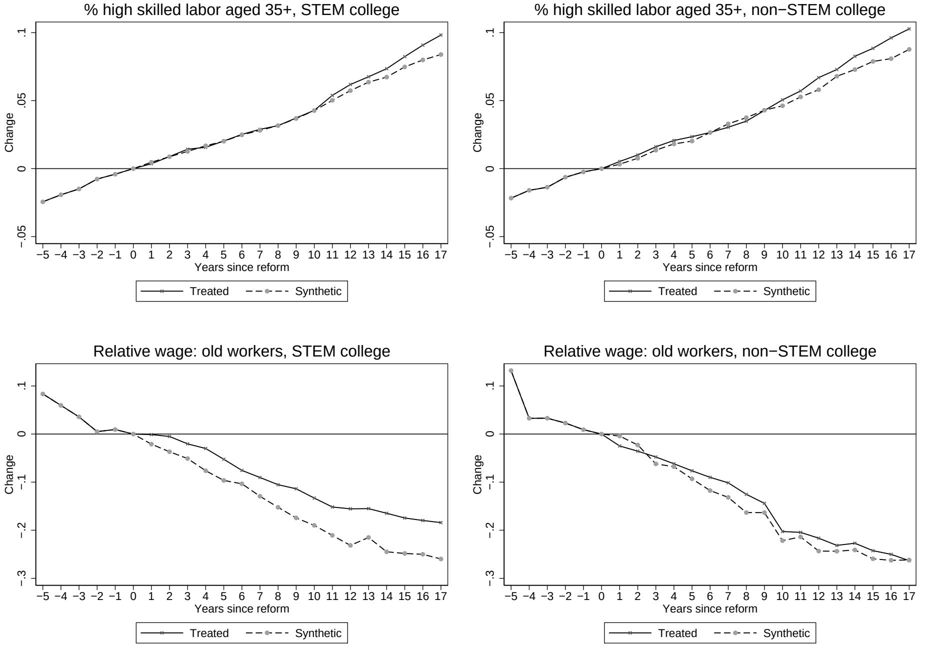

FIGURE 4. The Effects of the Reform on Old Workers: STEM vs Non-STEM Colleges. This figure presents the synthetic control estimates on skill composition and wages of old workers, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. Each 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.

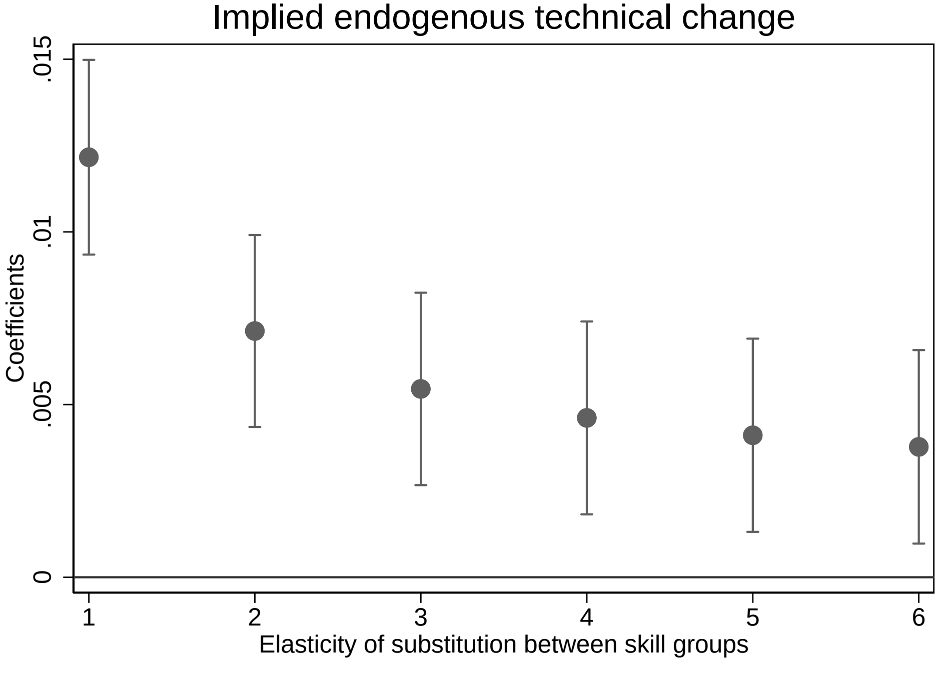

FIGURE 5. Implied Endogenous Technical Change for a Range of σ_E. This figure shows the estimated rate of endogenous SBTC (and the associated 95% C.I) from equation E.6 for given values of σ_E.

FIGURE 6. The Effects of the Reform on Investment in New Equipment. This figure presents the synthetic control estimates on investment in new equipment. 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.

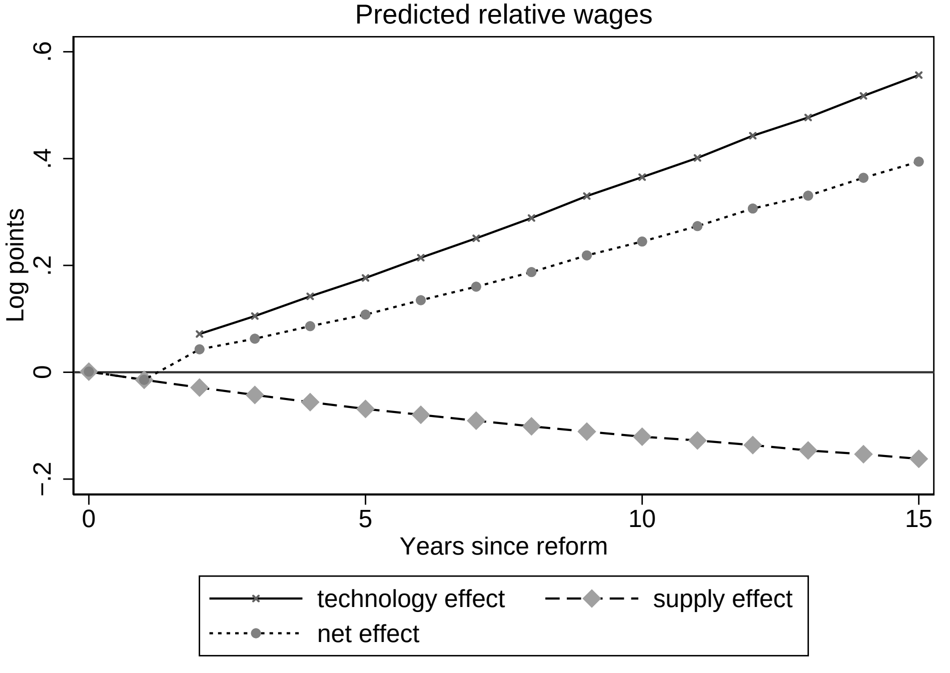

FIGURE 7. Predicted Relative Wages: Decomposing the Technology and Supply Effects. This figure reports the predicted relative wages from the production function estimates. See Section 4.3 for details.

TABLE 1. Comparison of Baseline Characteristics Before the Reform, by Treatment Status

Outcomes	Worker-level data Non-treated	Worker-level data Difference	Outcomes	Firm-level data Non-treated	Firm-level data Difference
Share of skilled workers, young	0.129***	0.002	Log output per worker	5.436***	-0.017
	(0.002)	(0.007)		(0.010)	(0.031)
Log high-skilled wages, young	12.387***	0.026	Plant size	48.202***	-1.602
	(0.006)	(0.018)		(1.670)	(5.203)
Log low-skilled wages, young	12.169***	-0.000	Hours worked	73195***	-1752
	(0.005)	(0.014)		(2499.6)	(7787.0)
Share of skilled workers, old	0.099***	0.021**	Log total wage costs	14.939***	0.018
	(0.003)	(0.009)		(0.020)	(0.061)
Log high-skilled wages, old	12.794***	0.012	Investment in machinery	809.656***	-115.901
	(0.004)	(0.013)		(31.803)	(99.067)
Log low-skilled wages, old	12.271***	0.012	Investment in machinery (exclu. transport)	667.045***	-68.377
	(0.007)	(0.021)		(28.932)	(90.125)
Growth in skill shares, young	0.002***	0.002	Investment in machinery and facilities	1146.9***	-21.370
	(0.001)	(0.002)		(47.869)	(149.116)
Growth in high-skilled wages, young	0.002	0.007	Employment by high-skilled industries (%)	0.483***	0.047
	(0.004)	(0.012)		(0.015)	(0.045)
Growth in low-skilled wages, young	-0.016***	0.001	Plants in high-skilled industries (%)	0.374***	0.016
	(0.001)	(0.003)		(0.010)	(0.031)
Growth in skill shares, old	0.002***	0.001	Outputs by high-skilled industries (%)	0.493***	0.051
	(0.000)	(0.001)		(0.015)	(0.047)
Growth in high-skilled wages, old	-0.015***	0.002	Log total costs of R and D	7.814***	-1.554**
	(0.002)	(0.006)		(0.260)	(0.718)
Growth in low-skilled wages, old	-0.014***	-0.001	Log total R and D man-years	3.336***	-1.621**
	(0.001)	(0.003)		(0.256)	(0.707)

Note: This table reports the mean pretreatment characteristics of the municipalities with treatment and the remaining municipalities without treatment. For growth rates outcomes, the table reports the annual change between 1967 and 1968. For the rest of the outcomes, the table reports the averages in 1967. All means are weighted by the number of plants in 1967.

TABLE 2. Estimates from the Relative Labor Demand Regression

	(1)	(2)	(3)	(4)
Aggr. Supply	-0.549	-0.546	-0.541	-0.492
	(0.348)	(0.341)	(0.325)	(0.312)
Trend	0.006	0.006	0.006	0.005
	(0.010)	(0.010)	(0.010)	(0.009)
Trend×Post	0.008**	0.009**	0.010**	0.011***
	(0.004)	(0.004)	(0.004)	(0.004)
Treated	0.015	0.021	0.027*	0.029**
	(0.018)	(0.016)	(0.015)	(0.014)
Post	-0.024	-0.038*	-0.057***	-0.075***
	(0.021)	(0.020)	(0.021)	(0.022)
Older worker	0.141***	0.141***	0.141***	0.141***
	(0.007)	(0.007)	(0.007)	(0.007)
Constant	-0.761	-0.756	-0.745	-0.644
	(0.713)	(0.698)	(0.665)	(0.639)
N	72	72	72	72

Note: This table reports the estimates of equations (E.5) and (E.7). See Section 3.3 and Appendix Section E for details. As independent variables, each regression includes a constant, a time trend (t), time trend interacted with the post dummy ($t \times P_{t(D)}$), a treatment group indicator (D), a post dummy ($P_{t(D)}$), an age group indicator (b_j), and the relative supply index. In column (1), we set M = 2. In columns (2) to (4), we set M = 3, 4, and 5, respectively.

TABLE 3. Production Function Estimates

Differences between treated and control group	CES (1)	CES (2)	Translog (3)	Translog (4)	Cobb-Douglas (5)	Cobb-Douglas (6)
Output elasticity: skilled labor						
2 years after reform	0.013	-0.001	0.011	-0.002	0.022	-0.002
	(0.035)	(0.029)	(0.038)	(0.027)	(0.030)	(0.024)
5 years after reform	0.042	0.019	0.044	0.031	0.044	0.018
	(0.032)	(0.027)	(0.041)	(0.032)	(0.029)	(0.023)
10 years after reform	0.090***	0.055*	0.100**	0.088*	0.080***	0.052*
	(0.031)	(0.032)	(0.049)	(0.052)	(0.029)	(0.030)
15 years after reform	0.138***	0.090**	0.156**	0.144*	0.116***	0.085**
	(0.037)	(0.045)	(0.061)	(0.077)	(0.032)	(0.042)
Average growth per year	0.010***	0.007**	0.011***	0.011**	0.007***	0.007**
	(0.003)	(0.003)	(0.003)	(0.005)	(0.002)	(0.003)
Output elasticity: unskilled labor						
2 years after reform	0.002	0.007	0.036	0.047	-0.007	0.007
	(0.034)	(0.030)	(0.033)	(0.033)	(0.029)	(0.025)
5 years after reform	-0.033	-0.017	-0.013	-0.003	-0.034	-0.016
	(0.030)	(0.026)	(0.034)	(0.024)	(0.028)	(0.022)
10 years after reform	-0.092***	-0.057*	-0.096*	-0.087	-0.080***	-0.054*
	(0.032)	(0.034)	(0.056)	(0.059)	(0.030)	(0.032)
15 years after reform	-0.150***	-0.098*	-0.179**	-0.172*	-0.126***	-0.093*
	(0.042)	(0.053)	(0.085)	(0.104)	(0.039)	(0.050)
Average growth per year	-0.012***	-0.008*	-0.016**	-0.016*	-0.009***	-0.008*
	(0.003)	(0.005)	(0.007)	(0.009)	(0.003)	(0.004)
Skill share control function	Yes	Yes	Yes	Yes	Yes	Yes
LP control function	No	Yes	No	Yes	No	Yes

Note: This table reports the predicted differences in output elasticity (by skill group) between the treated group and the control group at years, 2, 5, 10, and 15 after the reform (holding labor inputs at the mean). For the estimating production functions, see Appendix Section I. Columns (1), (3) and (5) include the control function for skill compositions and columns (2), (4) and (6) further control for unobserved productivity shocks via intermediate inputs. Average growth per year refers to the mean annual change in output elasticity starting from 2 years after the reform. Number of observations = 18441. Standard errors are clustered at the municipality level and given in parentheses.

TABLE 4. Production Function Estimates: Robustness Checks

Variable	(1)	(2)	(3)	(4)	(5)	(6)
Output elasticity: skilled labor						
Average growth per year	0.005**	0.006***	0.007***	0.006***	0.007***	0.006**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Output elasticity: unskilled labor						
Average growth per year	-0.007**	-0.009***	-0.009***	-0.008***	-0.010**	-0.011***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)
Output elasticity: capital						
Average growth per year					0.001	
					(0.004)	
Output elasticity: equipment capital						
Average growth per year						0.003
						(0.003)

Note: This table reports the predicted differences in the annual growth of output elasticity (starting from 2 years after the reform, by skill group) between the treated group and the control group (holding labor inputs at the mean). The estimating production function takes the form of Cobb-Douglas production function, under different assumptions of $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$ (equation (I.2) in Appendix Section I). 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 (including the share of skilled workers and employment size relative to total manufacturing employment in the municipality, both measured in 1960). Column (4) includes treatment group indicator variables (including a treatment group indicator and an indicator for implementing the reform early). Column (5) allows for the reform to impact the productivity of capital. Column (6) is the same as column (5), except that we replace total capital with equipment capital in the production function. 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**	0.833***	1.299**	0.944***	1.246*	1.590
	(0.362)	(0.228)	(0.552)	(0.338)	(0.677)	(1.115)
N	859	859	859	469	469	469
Covariates	No S×Y	Baseline	Munic’ trend	No S×Y	Baseline	Munic’ trend

Note: Estimates from equation (7) in the text. Coefficient on D_{ct} shows the effect of the reform. Standard errors are clustered at the municipality level. See Section 5 in the text for details.

ONLINE APPENDIX

Appendix A: Characteristics of the New Colleges

The size of each new college was also decided by the Parliament based on White Papers from the Government. The Ministry of Education forecasted the expected demand for study places in each region, based on the size of the cohorts in each region, the number of students graduating from high school and college in the previous years, and the expected demand for higher education. Given the small size of the pre-existing college-educated labor force in these areas, these new colleges had a large impact on the skill composition of local labor markets. This is illustrated in Appendix Figure A3 which shows, for each post-reform year, the total number of students enrolled in the seven regional colleges (which started operating between 1969 and 1971) as a fraction of the college educated labour market in these municipalities in 1970. Eight years after the opening of these colleges, the number of students enrolled (the flow of college workers into these labour markets) was as large as 25% of college workers in 1970 (the stock of college workers in 1970 in these locations).

The regional colleges provided education programs of shorter duration than those offered in traditional colleges. The two- and three-year programs covered most fields of study already available in the larger existing universities. In addition, new programs were developed with two- or three-year duration. For instance, all new regional colleges offered a program in business administration (at the time, business education was only available at one business school in Norway). Over half of the colleges offered programs in the STEM fields.

It is worth emphasizing that there was very little research output produced by these regional colleges in the period under study. Prior to the 1990s, virtually none of the academic staff in these colleges had a Ph.D. degree (Johnsen, 1999). Appendix Figure A4 shows R&D personnel and R&D expenditures among the regional colleges as a share of the total in the higher education sector. Up to 1990, the size of the R&D activities among the regional colleges was dwarfed by the R&D activity undertaken by the large research universities already established, with a R&D personnel share of 2 percent and expenditures of less than 5 percent of total university expenditure. In the 1990s (outside our period of study) some of these colleges are turned into universities, but the expenditure share in R&D remains small. Another measure of research activity is the number of patents generated by these institutions. Unfortunately we only have access to patents data by college or university starting in the mid 1990s. Nevertheless, all patenting activity is completely dominated by large research universities, even as late as the 1990s (see Appendix Figure A12). There is no notable patenting activity by the institutions forming the focus of our study.[68]

68. Data on patents was kindly provided by Professor Hans Hvide from the data set used in Hvide and Jones (2018). Almost all issued patents are in medicine and technology, and the two biggest universities

Appendix B: 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 $c = 1,\ldots,J+1$, and time periods $t = 1,\ldots,T$. 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 $t \in \{1,\ldots,T_0\}$.

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

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

where

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

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

$$Y_{ct}^0 = \delta_t + \boldsymbol{\theta_t Z_c} + \boldsymbol{\lambda_t \mu_c} + \varepsilon_{ct} \tag{B.3}$$

where δ_t is an unknown common factor with constant factor loadings across all municipalities, $\boldsymbol{Z_c}$ 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 (B.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 (B.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

dominate in medicine (Bergen and Oslo), and the technical university (Trondheim) dominates completely in patents in technology.

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 $\Sigma_{j=2}^{J+1} w_j^* Y_{jt}$ as an estimator for Y_{1t}^0, where w_j^* is a weight for each potential control municipality j such that

$$\begin{aligned} \Sigma_{j=2}^{J+1} w_j^* Y_{j\tau} = Y_{1\tau}, \forall \tau \in \{1, \dots, T_0\} \\ \Sigma_{j=2}^{J+1} w_j^* \boldsymbol{Z}_j = \boldsymbol{Z_1} \\ \Sigma_{j=2}^{J+1} w_j^* = 1, w_j^* \geq 0 \end{aligned}$$

The vector $\boldsymbol{W}^{*=(w_2^*, \dots, w_{J+1}^*)'}$ 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 $k_1 = k_2 = \cdots = k_{T_0 - 1} = 0$ 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 $\boldsymbol{X_1} = (Z_1'; \bar{Y}_1^{K_0}, \dots, \bar{Y}_1^{K_4})'$ be a $k \times 1$ vector of pre-intervention characteristics for the treated municipality.[69] Similarly, $\boldsymbol{X_0}$ is a $k \times J$ matrix that contains the same variables for the potential control municipalities, where the jth column of $\boldsymbol{X_0}$ is $(Z_j'; \bar{Y}_j^{K_0}, \dots, \bar{Y}_j^{K_4})'$. The vector $\boldsymbol{W^*}$ is chosen to minimize the distance between X_1 and $X_0 W$, where

$$\|\boldsymbol{X_1} - \boldsymbol{X_0 W}\| = \sqrt{(\boldsymbol{X_1} - \boldsymbol{X_0 W})' \boldsymbol{V} (\boldsymbol{X_1} - \boldsymbol{X_0 W})} \tag{B.4}$$

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

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

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

69. In vector $\boldsymbol{Z_j}$, 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.

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 period as

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

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.[70] 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 $p_{\tau} = \frac{\sum_{l=1}^{K} \mathbf{1}(\widehat{\alpha}_{\tau}^{l} > \widehat{\alpha}_{\tau}^{0})}{K}$.[71]

Appendix C: 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.

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

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

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

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

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.[72] 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 πs should be close to zero and statistically insignificant.

In Appendix Tables A2 and A3, 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 (C.1). We allow the effects of the reform (the estimated τs) 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

72. The 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.

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.

Appendix D: 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 4).

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.[73] 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

73. Military 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.

residence:

$$Q_i = \chi_m + \chi_c + \pi R_{mc} + \varepsilon_{imc} \tag{D.1}$$

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 (D.1) by skill type 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 A10 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.

Appendix E: Estimating the Card and Lemieux (2001) Model: Details

We follow Card and Lemieux (2001) and estimate equation (2) in two steps, using the data generated in Figure 1.[74] In the first step, σ_A (the gross elasticity of substitution between different age groups j within a given skill group) is estimated from a regression of age-group- specific relative wages on age-group-specific relative supplies, age effects, and time effects.

$$\log \frac{w^s_{jt(D)}}{w^u_{jt(D)}} = b_j + \gamma_{tD} - \frac{1}{\sigma_A}(\log \frac{S_{jt(D)}}{U_{jt(D)}}) + e_{jtD} \tag{E.1}$$

where b_j and γ_{tD} 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. We explore exogenous changes in supply (and, therefore, identify the demand curve) by using

74. We 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.

differences within age (j) and across treatment groups (D) to estimate the following equation:[75]

$$\log \frac{w^s_{jt(1)}}{w^u_{jt(1)}} - \log \frac{w^s_{jt(0)}}{w^u_{jt(0)}} = (\gamma_{t1} - \gamma_{t0}) - \frac{1}{\sigma_A}(\log \frac{S_{jt(1)}}{U_{jt(1)}} - \log \frac{S_{jt(0)}}{U_{jt(0)}}) + (e_{jt1} - e_{jt0}) \quad \text{(E.2)}$$

Under the assumption of this model, σ_A is identified from exogenous shifts in age specific skill supplies, since technology is assumed to operate uniformly across age groups (and therefore it can be subsumed in γ_{tD}).

Given the estimate of σ_A, the efficiency parameters β^s_j and β^u_j (which are assumed to be invariant to D) are estimated using the following equations (as in Card and Lemieux (2001)):

$$\log w^s_{jt(D)} + \frac{1}{\sigma_A} \log S_{jt(D)} = \gamma^s_{tD} + \log(\beta^s_j) + e^s_{jtD} \quad \text{(E.3)}$$

$$\log w^u_{jt(D)} + \frac{1}{\sigma_A} \log U_{jt(D)} = \gamma^u_{tD} + \log(\beta^u_j) + e^u_{jtD} \quad \text{(E.4)}$$

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

In the second step of the estimation, we use data from both the treated group and the control group to identify the effects of college openings on technology change. It is here that we face the challenge of separately identifying σ_E and endogenous technical change parameters. Based on equation (2), and equipped with the estimated σ_A and aggregate supplies, we can rewrite the model as:

$$\log \frac{w^s_{jt(D)}}{w^u_{jt(D)}} + \frac{1}{\sigma_A}(\log \frac{S_{jt(D)}}{U_{jt(D)}} - \log \frac{S_{t(D)}}{U_{t(D)}}) = \delta_{t(D)} + b_j - \frac{1}{\sigma_E} \log \frac{S_{t(D)}}{U_{t(D)}} + e_{jtD} \quad \text{(E.5)}$$

where b_j are age-group dummies, and $\delta_{t(D)} \equiv \log(\frac{\theta_{st(D)}}{\theta_{ut(D)}})$ represents the relative technology efficiency which is specific to each year and treatment group. Our parameters of interest are the sequence of $\delta_{t(D)}$ and σ_E.

In this equation, $\frac{S_{t(D)}}{U_{t(D)}}$ could be correlated with e_{jtD}, leading to biased estimates of all the parameters. More importantly, any exogenous variation in skill supplies is no

75. In the production function, the form of technical change is common to both age groups and there is no age-specific technical change: β^s_j and β^u_j do not vary with t. With this assumption, σ_A can be identified in the data using within municipality and skill movements in age-specific supplies as we have shown. Note that our approach can still identify σ_A from the data even if we allow for age-specific technical change, as long as the age-specific technical change is exogenous (i.e., if they evolve over time in the same way between the treatment group and the control group).

longer a valid exclusion restriction for $\frac{S_{t(D)}}{U_{t(D)}}$ in this equation because it also has a direct effect on technical change, $\delta_{t(D)}$, through the channel we emphasize in the paper.

To address this issue, we proceed with two different identification strategies. One is to use an external estimate of σ_E from the literature to back out the technical change parameters. An advantage of this approach is that we can experiment with a range of plausible values of σ_E to gauge the amount of technical change that is needed to match our data. Given an external estimate of σ_E, our regression model becomes:

$$\log \frac{w^s_{jt(D)}}{w^u_{jt(D)}} + \frac{1}{\sigma_A}(\log \frac{S_{jt(D)}}{U_{jt(D)}} - \log \frac{S_{t(D)}}{U_{t(D)}}) + \frac{1}{\sigma_E} \log \frac{S_{t(D)}}{U_{t(D)}} = \delta_{t(D)} + b_j + e_{jtD} \qquad \text{(E.6)}$$

Following much of the literature on this topic, the relative technology efficiency, we model $\delta_{t(D)}$ as a linear trend (although we could relax this assumption), interpreted as skill-biased technical change (Katz and Murphy, 1992):[76] $\delta_{t(D)} = \delta_0 t + \delta_1 (t \times D) + \delta_2 D$. The trend in technical change is allowed to vary with treatment. δ_0 represents skill-biased technical change in the synthetic control group, whereas δ_1 represents the incremental skill-biased technical change taking place in the treated group. A positive δ_1 implies endogenous technical change.

A second idea is to make one additional timing assumption in order to identify both σ_E and $\delta_{t(D)}$ from our data. A reasonable possibility is to assume that $\frac{\theta_{st(D)}}{\theta_{ut(D)}}$ does not evolve differentially with D in the years immediately following the reform. Formally, 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 in equation (E.5) for fixed $\frac{\theta_{st(D)}}{\theta_{ut(D)}}$, by relating differences in relative wages to (exogenous) 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 $\frac{\theta_{st(D)}}{\theta_{ut(D)}}$, $t > M$.[77]

76. The 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 M 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 M years after the opening of the college).

77. Another intuitive idea would be to use the older workers in the years immediately after the reform to identify $\log(\frac{\theta_{st(1)}}{\theta_{ut(1)}}) - \log(\frac{\theta_{st(0)}}{\theta_{ut(0)}})$ because they did not experience increases in $\frac{S_{jt(D)}}{U_{jt(D)}}$ until much later. Given $\log(\frac{\theta_{st(1)}}{\theta_{ut(1)}}) - \log(\frac{\theta_{st(0)}}{\theta_{ut(0)}})$, 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.

Under the additional timing assumption, we estimate equation (E.5) by modeling $\delta_{t(D)}$ as

$$\delta_{t(D)} = \hat{\delta}_0 t + \hat{\delta}_1 (t \times P_{t(D)}) + \hat{\delta}_2 P_{t(D)} + \hat{\delta}_3 D, \tag{E.7}$$

where $P_{t(D)}$ is an indicator function that takes value 1 if $t \geq M$ and $D = 1$ (i.e., if the observation is after M years since the reform and is in the treated group) and takes value zero elsewhere, and $\hat{\delta}_1$ represents the incremental skill-biased technical change taking place in the treated group. A positive $\hat{\delta}_1$ implies endogenous technical change. Our benchmark case is $M = 2$. Below, we examine the robustness of our estimates to $M = \{3, 4, 5\}$.

E.1. Allowing σ_A to Differ by Skill Groups

In the model developed in Section 3.3, we have assumed that the elasticity of substitution across age groups the same for both skill groups. The model is easily generalized by introducing separate elasticities of substitution, σ_{As}, and σ_{Au}, for skilled and unskilled workers, respectively. Under the assumption that wages are equated to marginal products, we can derive a pair of wage determination equations:

$$\log w^s_{jt(D)} = \log(\theta_{st(D)}) + \log \beta_j + (\frac{1}{\sigma_{As}} - \frac{1}{\sigma_E}) \log \tilde{S}_{t(D)} - \frac{1}{\sigma_{As}} (\log S_{jt(D)}) \tag{E.8}$$

$$\log w^u_{jt(D)} = \log(\theta_{ut(D)}) + \log \alpha_j + (\frac{1}{\sigma_{Au}} - \frac{1}{\sigma_E}) \log \tilde{U}_{t(D)} - \frac{1}{\sigma_{Au}} (\log U_{jt(D)}) \tag{E.9}$$

where $\tilde{S}_{t(D)}$ and $\tilde{U}_{t(D)}$ are the same labor aggregates as defined previously, except that the elasticity of substitution is now specific to each education group. We estimate σ_{As} and σ_{Au} separately, by exploring exogenous changes in absolute supply within age (j) and across treatment groups (D):

$$\log w^s_{jt(1)} - \log w^s_{jt(0)} = (d^s_{t1} - d^s_{t0}) - \frac{1}{\sigma_{As}} (\log S_{jt(1)} - \log S_{jt(0)}) + \left(e^s_{jt1} - e^s_{jt0}\right) \tag{E.10}$$

$$\log w^u_{jt(1)} - \log w^u_{jt(0)} = (d^u_{t1} - d^u_{t0}) - \frac{1}{\sigma_{Au}} (\log U_{jt(1)} - \log U_{jt(0)}) + \left(e^u_{jt1} - e^u_{jt0}\right) \tag{E.11}$$

Therefore, the separate elasticities of substitution, σ_{As} and σ_{Au}, can be identified from regressions of changes in absolute wages on changes in absolute supply, controlling for year fixed effects.

Given the estimate of σ_{Au} and σ_{As}, the efficiency parameters β^s_j and β^u_j are estimated via the following equations similar to equations E.3 and E.4. With the estimated σ_{As}, σ_{Au}, $\log(\beta^s_j)$ and $\log(\beta^u_j)$, we construct estimates of the aggregate supplies of skilled and unskilled labor in each year for both the treated and synthetic groups. We use data from both the treated group and the synthetic group to estimate the following regression of relative wages:

$$\log \frac{w^s_{jt(D)}}{w^u_{jt(D)}} - (\log \frac{\tilde{S}^{1/\sigma_{As}}_{t(D)}}{\tilde{U}^{1/\sigma_{Au}}_{t(D)}} \frac{S^{-1/\sigma_{As}}_{jt(D)}}{U^{-1/\sigma_{Au}}_{jt(D)}}) = \delta_{t(D)} + b_j - \frac{1}{\sigma_E} \log \frac{\tilde{S}_{t(D)}}{\tilde{U}_{t(D)}} + e_{jtD} \qquad \text{(E.12)}$$

where $\delta_{t(D)} \equiv \log(\frac{\theta_{st(D)}}{\theta_{ut(D)}})$, which represents the relative technology efficiency that is specific to the year and treatment group.

Columns (2) and (3) in Appendix Table A5 report the estimated σ_{As}, and σ_{Au}, respectively. We find that the elasticity of substitution across age groups differ by skill groups: the implied σ_{As} is 3 (=1/0.329), which is significantly different from the estimated σ_{Au} (=1/0.196=5). However, allowing for separate elasticities of substitution between age groups does not affect the implied skill-biased technical change. Appendix Table A6 shows that the implied skill-biased technical change is virtually unchanged relative to our baseline specification where σ_{As} and σ_{Au} is forced to be the same.

Appendix F: A Theoretical Model of Endogenous Technology Adoption

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.

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:

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

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):

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

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.

Appendix G: Supply Effects and Technical Change in the National Data

In this section, we follow Card and Lemieux (2001) and estimate supply and technology effects using the national data from Norway. From the worker-level data described in Section 2.2 in the paper, we construct aggregate time series of relative supply and relative wage, by age groups. In the first step, σ_A (the gross elasticity of substitution between different age groups j within a given skill group) is estimated from a regression of age-group- specific relative wages on age-group-specific relative supplies, age effects, and time effects.

$$\log \frac{w^s_{jt}}{w^u_{jt}} = b_j + \gamma_t - \frac{1}{\sigma_A}(\log \frac{S_{jt}}{U_{jt}}) + e_{jt} \tag{G.1}$$

where b_j and γ_t are indicators for age and year effects, respectively. The year effects absorb both the relative technology efficiency between skilled and unskilled labor, as well as any effects of changing aggregate supply. Column (1) of Appendix Table A7 shows the results from this regression. The implied σ_A is 2.28.

Given the estimate of σ_A, the efficiency parameters β^s_j and β^u_j are estimated using the following equations (as in Card and Lemieux (2001)):

$$\log w^s_{jt} + \frac{1}{\sigma_A}\log S_{jt} = \gamma^s_t + \log(\beta^s_j) + e^s_{jt} \tag{G.2}$$

$$\log w^u_{jt} + \frac{1}{\sigma_A}\log U_{jt} = \gamma^u_t + \log(\beta^u_j) + e^u_{jt} \tag{G.3}$$

These equations are derived by equalizing the marginal product of labor with the wage for each combination of age and skill groups. γ^s_t and γ^u_t is a set of year dummies, and $\log(\beta^s_j)$ and $\log(\beta^u_j)$ are estimated from the age effects in the above equations. With the estimated σ_A, $\log(\beta^s_j)$ and $\log(\beta^u_j)$, we construct estimates of the aggregate supplies of skilled and unskilled labor in each year.

In the second step of the estimation, if we ignore endogenous technical change and assume that exogenous technical change follows a linear trend, we can estimate the following regression model:

$$\log \frac{w^s_{jt}}{w^u_{jt}} + \frac{1}{\sigma_A}(\log \frac{S_{jt}}{U_{jt}} - \log \frac{S_t}{U_t}) = \delta t + b_j - \frac{1}{\sigma_E}\log \frac{S_t}{U_t} + e_{jt} \tag{G.4}$$

where σ_A is estimated from the first-step regression, b_j are age-group dummies, σ_E is the elasticity of substitution between skilled and unskilled labor, and δ represents the rate of exogenous technical change. In column (2) of Appendix Table A7, we report the results from this regression. The implied σ_E is 1.32, and the rate of exogenous technical change is 0.021 per annum.

When there is endogenous technical change, then $\frac{S_t}{U_t}$ is correlated with e_{jt} and the estimated parameters σ_E and δ are biased. To understand the extent of the bias, we use the implied rate of endogenous technical change identified from the treatment-control data in Section 3.3. Let θ be the implied rate of endogenous technical change per unit change of relative supply.[78] The total contribution of endogenous technical change is

78. The differential technical change between treated and control group is 0.008 per year (first column in Table 2). From year 2 to year 17, the implied treatment-control difference is 0.008×15=0.12 and the

$\theta \times (\log \frac{S_t}{U_t} - \log \frac{S_0}{U_0})$, where $\frac{S_0}{U_0}$ is the initial relative skill level in the first year in the data (year 1967). Subtracting the contribution from endogenous technical change, we can identify σ_E and the rate of exogenous technical change from the regression model:

$$\log \frac{w^s_{jt}}{w^u_{jt}} + \frac{1}{\sigma_A}(\log \frac{S_{jt}}{U_{jt}} - \log \frac{S_t}{U_t}) - \theta \times (\log \frac{S_t}{U_t} - \log \frac{S_0}{U_0}) = \hat{\delta} t + \hat{b_j} - \frac{1}{\hat{\sigma_E}} \log \frac{S_t}{U_t} + \hat{e_{jt}} \tag{G.5}$$

Column (3) of Appendix Table A7 shows the results from this regression. The implied σ_E is less than 1 and much smaller than before, which indicates that ignoring endogenous technical change leads to significant bias in the estimation of σ_E. The estimated rate of exogenous technical change remains unaffected, due to the assumption that endogenous technical change is linear in aggregate relative supply. In Appendix Figure A11, we show the cumulative contribution of exogenous and total technical change to relative wages since year 1967 (1967 is normalized to 0). We find that incorporating our estimate of endogenous technical change in the empirical model more than doubles the total skill-biased technical change taking place between 1967 and 1990.

Appendix H: 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}. Notice that:

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

where

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

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 across different t, Q_{jct} remains fixed at its initial value in 1960 (taken from the 1960 Census). This implies that the share of skilled workers in a given industry within a municipality is a constant proportion of the share of skilled workers in that

treatment-control gap in relative aggregate supply in year 17 is 0.17. Therefore, the endogenous technical change per unit change of relative supply, θ, is equal to 0.12/0.17.

municipality. Specifically, we write:

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

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

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 I: Production Function: Estimation Details

As discussed in Section 4, when estimating the production function, we face two potential endogeneity problems: one is L_{jct} (quantity of labor employed) and K_{jct} being endogenous and the other is π_{ct} (skill composition of employment) being endogenous. In this section, we begin by discussing the details of implementing the control function approach to take into account of the endogeneity of skill composition of employment (assuming K_{jct} and L_{jct} are exogenous). We then specify the second control function that takes into account the endogeneity of quantity of labor employed and capital.

I.1. First Stage: Estimating the Skill-share Equation

We first regress skill shares in municipality c and year t on reform indicators and a set of control variables:

$$\pi_{ct(D)} = \theta_1 D_{ct} + \theta_2 D_{ct} \times t_D + \gamma_c^s + \gamma_{k(c)t}^s + v_{ct} \tag{I.1}$$

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 $t_D > 0$.

We recover $\widehat{v}_{ct}$, the estimated residuals from this regression, which we include as regressors when estimating the production functions.

I.2. Second Stage: Estimating the Production Functions

Next, we estimate the parameters of the production functions, including the estimated residuals of the first-stage regression as additional regressors, as specified in equation (6). Our first control function rests on the assumption that the disturbances in the production functions are independent of the exclusion restriction–variables that affect the output only through the skill compositions (but not through technology change). We assume that D_{ct} and $D_{ct} \times t_D$ in the first-stage regression are exogenous (conditional on covariates), and, therefore lead to exogenous variation in π_{ct}. However, they are not valid exclusion restrictions because they could affect output via technical change. Following the discussion in Section 4, as exclusion restriction, we exploit the timing of the college expansion reform. The identifying assumption is that technologies do not respond to the changes in the skill composition π_{ct} in the first M years immediately following the opening of a college.

We estimate three different specifications of production functions: the Cobb-Douglas, the CES, and the translog. The advantage of the translog specification is that the elasticity of substitution between different skill types is not a constant. The advantage of the Cobb-Douglas specification is that it is much simpler than the linearized CES or the translog described above, and the impacts of the reform on factor neutral productivity, and the productivity of skilled and unskilled labour can be read directly from the estimated coefficients because output is log linear in inputs. The disadvantage is that the Cobb Douglas production function restricts the elasticity of substitution between skilled and unskilled labor to one. If it is larger than one, then an increase in the supply of skill may increase the share of the wage bill going to skilled workers, which may be all we are capturing with the Cobb Douglas specification. In Appendix Section J, we provide a numerical exercise in which we show that an increase in factor supplies has limited impact on factor shares under a reasonable range of elasticity of substitution parameters between skilled and unskilled labor.[79]

Cobb-Douglas production function

For the Cobb-Douglas production function, we estimate the following second-stage regression:

$$\log(Y_{jct}) = \beta_{0,jct(D)} + \alpha \log K_{jct} + \beta_{1,ct(D)} \log S_{jct(D)} + \beta_{2,ct(D)} \log U_{jct(D)} + \rho_1 \widehat{v}_{ct} + \rho_2 \widehat{v}_{ct}^2 + \mu_{jct} \tag{I.2}$$

CES production function

For the CES production function, we follow Kmenta (1967) to linearize the CES aggregate (equation (1)) around $\rho = 0$ using a second-order Taylor expansion. We

79. The Cobb–Douglas production function is a special case of the CES production function discussed in Section 3.3 (with capital), where $\rho = 0$. $\beta_{1,ct(D)}$ and $\beta_{2,ct(D)}$ correspond to the extensive technical change parameters in the CES production function (see Section 3.3 for details), although we do not restrict the production function to be constant return to scale. The TFP parameter, $\beta_{0,jct(D)}$, is a weighted average of the labor-augmenting technical change parameters in the CES function ($\beta_{0,jct(D)} = \alpha_t(D)a_{t(D)} + (1 - \alpha_t(D))b_{t(D)}$).

estimate the following second-stage regression:

$$\log(Y_{jct}) = \beta_{0,jct(D)} + \alpha \log K_{jct} + \beta_{1,ct(D)} \log S_{jct(D)} + \beta_{2,ct(D)} \log U_{jct(D)} + \beta_{3,ct(D)} (\log \frac{S_{jct(D)}}{U_{jct(D)}})^2 + \rho_1 \widehat{v}_{ct} + \rho_2 \widehat{v}_{ct}^2 + \mu_{jct} \tag{I.3}$$

Translog production function

For the translog production function, we estimate the following second-stage regression:

$$\begin{aligned}\log(Y_{jct}) = \beta_{0,jct(D)} &+ \alpha \log K_{jct} + \beta_{1,ct(D)} \log S_{jct(D)} + \beta_{2,ct(D)} \log U_{jct(D)} \\ &+ \beta_{3,ct(D)} (\log S_{jct(D)})^2 + \beta_{4,ct(D)} (\log U_{jct(D)})^2 + \beta_{5,ct(D)} \log U_{jct(D)} \log S_{jct(D)} \\ &+ \rho_1 \widehat{v}_{ct} + \rho_2 \widehat{v}_{ct}^2 + \mu_{jct}\end{aligned} \tag{I.4}$$

In each specification of the production function, the reform indicator, D, can affect the technology of production via the labor productivity parameters ($\beta_{\kappa,ct(D)}, \kappa \geq 1$) and factor-neutral productivity ($\beta_{0,jct(D)}$). In our baseline results (reported in Table 3), these productivity parameters depend on industry, municipality, time, and lags of the reform according to the following model:

$$\beta_{0,jct(D)} = \gamma_{jc}^0 + \gamma_{jt}^0 + \gamma_{k(c)t}^0 + \delta_1^0 (t_D \times P_{t(D)}) + \delta_2^0 P_{t(D)} \tag{I.5}$$

$$\beta_{\kappa,ct(D)} = \gamma_t^\kappa + \delta_1^\kappa (t_D \times P_{t(D)}) + \delta_2^\kappa P_{t(D)}, \forall \kappa \geq 1 \tag{I.6}$$

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). γ_{jc}^0 are industry by municipality fixed effects, γ_{jt}^0 are industry-by-year effects, $\gamma_{k(c)t}^0$ are county-by-year effects, and γ_t^κ are year effects.[80] $P_{t(D)}$ is an indicator function that takes value 1 if $t_D \geq M$, where $M = 2$. This parameterization allows the output elasticities of skilled and unskilled labor to vary with time, location, and with the reform.[81]

We also experiment with alternative specifications, where, in addition to time and the reform, the output elasticities of skilled and unskilled labor are allowed to vary

80. Industry-year fixed effects fully absorb any permanent heterogeneity at the industry and municipality level. Industry-year fixed effects absorb any industry-specific time-varying shocks at the national level. 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.

81. We 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 (see column 5 of Table 4). 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. We have also estimated the same model as column (5) but replace capital with equipment capital. Column 6 shows that the results do not materially change. 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.

with time, location, industry and with the reform. For instance, we expand equation (I.6) by including industry fixed effects (column 2 in Table 4). We also include initial municipality-industry specific characteristics and treatment group indicators in equation (I.6) (columns 3 and 4 in Table 4).[82] Finally, we allow for the productivity of capital (α) to vary by time and the reform in the same way as the labor productivity parameters (column 4 in Table 4), where:

$$\alpha_{ct(D)} = \gamma_t^K + \delta_1^K (t_D \times P_{t(D)}) + \delta_2^K P_{t(D)}$$

Regressions are weighted by the number of plants in each municipality-industry cell (fixed at the levels in 1967, the first year in the panel). Standard errors are clustered at the municipality level.

I.3. Additional control function using Levinsohn and Petrin (2003)

To address the endogeneity problem of labor and capital inputs, we also estimate another specification of the production function, adding the terms in equation (I.7) as controls (in addition to polynomials in $\widehat{v}_{jct}$). Levinsohn and Petrin (2003) uses a structural model of an optimizing firm to derive the conditions under which the intermediate input demand function depends on the firm-specific state variables, including productivity shocks and capital.

Following Levinsohn and Petrin (2003), suppose that the demand for intermediate input (material inputs in our case), m_{jct}, depends on the firm's stock of capital and (factor-neutral) productivity shock, ω_{jct}. Under the assumption that the intermediate input demand is monotonic in productivity shock for all capital, the intermediate input demand function can be inverted to yield a control function for the unobserved productivity shock. Specifically, LP shows that ω_{jct}, the unobserved productivity shock, can be written as a function of $\log K_{jct}$ and $\log m_{jct}$:

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

Let $\varphi_t(\log m_{jct}, \log K_{jct}) = \alpha \log K_{jct} + \omega_t(m_{jct}, K_{jct})$. LP use a third-order polynomial to approximate $\varphi_t(\log m_{jct}, \log K_{jct})$ nonparametrically:

$$\varphi(\log m_{jct}, \log K_{jct}) = \Sigma_{p=0}^{3} \Sigma_{q=0}^{3-p} \delta_{pq} \log K_{jct}^p \log m_{jct}^q \tag{I.7}$$

82. As the initial municipality-industry specific characteristics, we use the level of the share of skilled workers and the share of employment (within manufacturing sectors) working for an industry in a given municipality. Both of these variables are constructed from the census data in 1960 (a pre-reform year). As treatment group indicators, we include a treatment group indicator and another “early-reformer" indicator based on the timing of the treatment.

Adding the second control function, the estimating equation for the Cobb-Douglas production function (as an example) becomes

$$\log(Y_{jct}) = \beta_{0,jct(D)} + \Sigma_{p=0}^{3}\Sigma_{q=0}^{3-p}\delta_{pq}\log K_{jct}^{p}\log m_{jct}^{q} + \beta_{1,ct(D)}\log S_{jct(D)} + \beta_{2,ct(D)}\log U_{jct(D)}$$
$$+ \beta_{3,ct(D)}(\log\frac{S_{jct(D)}}{U_{jct(D)}})^2 + \rho_1\widehat{v}_{ct} + \rho_2\widehat{v_{ct}^2} + \mu_{jct}$$

where m_{jct} is the total material inputs used in industry j, municipality c and year t. The additional control function, $\Sigma_{p=0}^{3}\Sigma_{q=0}^{3-p}\delta_{pq}\log K_{jct}^{p}\log m_{jct}^{q}$ is a flexible function of capital and intermediate inputs, both of which are observed in the data. Conditional on the control functions, the remaining error term in the production function (μ_{jct}) is assumed not correlated with total employment, capital and skill compositions of employment.[83]

Appendix J: Numerical Exercise: The Impacts of Technical Change and Factor Supplies on Factor Shares

Ignore the age dimension (as we do in the production function section), and take a standard CES:

$$Y_{ct} = \left(\theta_{St}S_{ct}^{\rho} + \theta_{Ut}U_{ct}^{\rho}\right)^{\frac{1}{\rho}}$$

Then work out wages and factor shares:

$$w_{Sct} = \left(\theta_{St}S_{ct}^{\rho} + \theta_{Ut}U_{ct}^{\rho}\right)^{\frac{1}{\rho}-1}\theta_{St}S_{ct}^{\rho-1}$$

$$w_{Sct}S_{ct} = \left(\theta_{St}S_{ct}^{\rho} + \theta_{Ut}U_{ct}^{\rho}\right)^{\frac{1}{\rho}-1}\theta_{St}S_{ct}^{\rho}$$

$$SH_{Sct} = \frac{w_{Sct}S_{ct}}{Y_{ct}} = \frac{\theta_{St}S_{ct}^{\rho}}{\left(\theta_{St}S_{ct}^{\rho} + \theta_{Ut}U_{ct}^{\rho}\right)}$$

$$SH_{Uct} = \frac{w_{Sct}U_{ct}}{Y_{ct}} = \frac{\theta_{Ut}U_{ct}^{\rho}}{\left(\theta_{St}S_{ct}^{\rho} + \theta_{Ut}U_{ct}^{\rho}\right)}$$

Let $R_{ct} = \frac{S_{ct}}{U_{ct}}$. Then we can rewrite these as:

$$SH_{Sct} = \frac{w_{Sct}R_{ct}}{Y_{ct}} = \frac{\theta_{St}R_{ct}^{\rho}}{\left(\theta_{St}R_{ct}^{\rho} + \theta_{Ut}\right)}$$

83. The 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.

Finally:[84]

$$D_{ct} = \frac{\partial SH_{Sct}}{\partial R_{ct}} = \frac{\theta_{St}\rho R_{ct}^{\rho-1}\left(\theta_{St}R_{ct}^{\rho}+\theta_{Ut}\right) - \theta_{St}R_{ct}^{\rho}\theta_{St}\rho R_{ct}^{\rho-1}}{\left(\theta_{St}R_{ct}^{\rho}+\theta_{Ut}\right)^2} = \frac{\theta_{St}\rho R_{ct}^{\rho-1}\theta_{Ut}}{\left(\theta_{St}R_{ct}^{\rho}+\theta_{Ut}\right)^2} =$$

$$= \frac{\rho}{R_{ct}}\frac{\theta_{St}R_{ct}^{\rho}}{\left(\theta_{St}R_{ct}^{\rho}+\theta_{Ut}\right)}\frac{\theta_{Ut}}{\left(\theta_{St}R_{ct}^{\rho}+\theta_{Ut}\right)} = \frac{\rho}{R_{ct}}SH_{Sct}SH_{Uct}$$

In the Cobb-Douglas production, we estimate:

$$Y_{ct} = \theta_t S_{ct}^{\theta_{St}} U_{ct}^{\theta_{Ut}} K_{ct}^{\theta_{Kt}} \varepsilon_{ct}$$

We estimate θ_{St} and θ_{Ut} so we can compute:

$$SH'_{Sct} = \frac{\theta_{St}}{\theta_{St}+\theta_{Ut}}$$

Let $t=0$ be the year just before the reform. We can generate two series:

$$SH'_{Sct} = \frac{\theta_{St}}{\theta_{St}+\theta_{Ut}}$$

and SH'^{*}_{Sct} and D'_{Sct} such that

$$D'_{Sct} = \frac{\partial SH'_{Sct}}{\partial R_{ct}} = \frac{\rho}{R_{ct}}SH'_{Sct}SH'_{Uct}$$

$$SH'^{*}_{Sc0} = SH'_{Sc0}$$

$$SH'^{*}_{Sc1} = SH'^{*}_{Sc0} + D'^{*}_{Sc0}(R_{c1}-R_{c0})$$

$$SH'^{*}_{Sct} = SH'^{*}_{Sct-1} + D'^{*}_{Sct-1}(R_{ct}-R_{ct-1})$$

$$D'^{*}_{Sct} = \frac{\partial SH'^{*}_{Sct}}{\partial R_{ct}} = \frac{\rho}{R_{ct}}SH'^{*}_{Sct}SH'^{*}_{Uct}$$

84. We can add capital to the model. In that case, we would look at the shares of high and low skill pay on the total wage bill. Suppose we had the following production function:

$$Y_{ct} = \left(\theta_{Lt}L_{ct}^{\varphi}+\theta_{Kt}K_{ct}^{\varphi}\right)^{\frac{1}{\varphi}}$$

where

$$L_{ct} = \left(\theta_{St}S_{ct}^{\rho}+\theta_{Ut}U_{ct}^{\rho}\right)^{\frac{1}{\rho}}.$$

One can show that the share of high and low skill pay on the total wage bill has the exact form as the one we derive below.

SH'^{*}_{Sct} can be calculated for different values of ρ, so we will label it accordingly: $SH'^{*}_{Sct}(\rho)$. We will do it mainly for $\rho > 0$ (and $\rho < 1$), which is the case most consistent with the literature.

In Appendix Table A14, we report compare SH'_{Sct} and SH'^{*}_{Sct} at two values of ρ. SH'_{Sct} tells us about the actual change in factor shares as implied by the Cobb-Douglas estimates, and which can incorporate genuine technical change from automatic increase in factor shares because of an increase in factor supplies (and elasticity of substitution between factors larger/different than one). $SH'^{*}_{Sct}(\rho)$ tells us, for a given ρ, what the change in the share would be in the absence of technical change (solely due to changes in R_{ct}). For reasonable values of $\rho = 0.5$, the two series are very different (compare column (2) to column (1)), and they only become similar when ρ gets unrealistically close to 1 (compare column (3) to column (1)).

APPENDIX TABLES AND FIGURES

FIGURE A1. Completed Level of Education for Individuals Aged 16+ (in percentages): 1960–1990

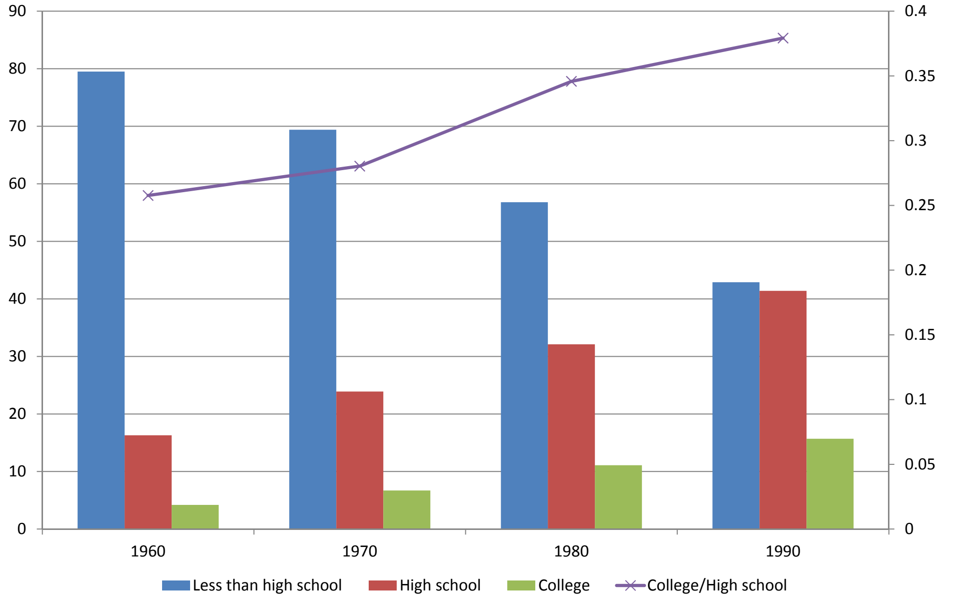

This figure shows the distribution of levels of completed education in the adult population (in percentages; left scale) by year. The line shows the ratio of the college-educated persons over high-school educated persons by year (right scale). Source: Statistics Norway (1994).

FIGURE A2. Geographic Locations of New Colleges in Norway: 1967–1985

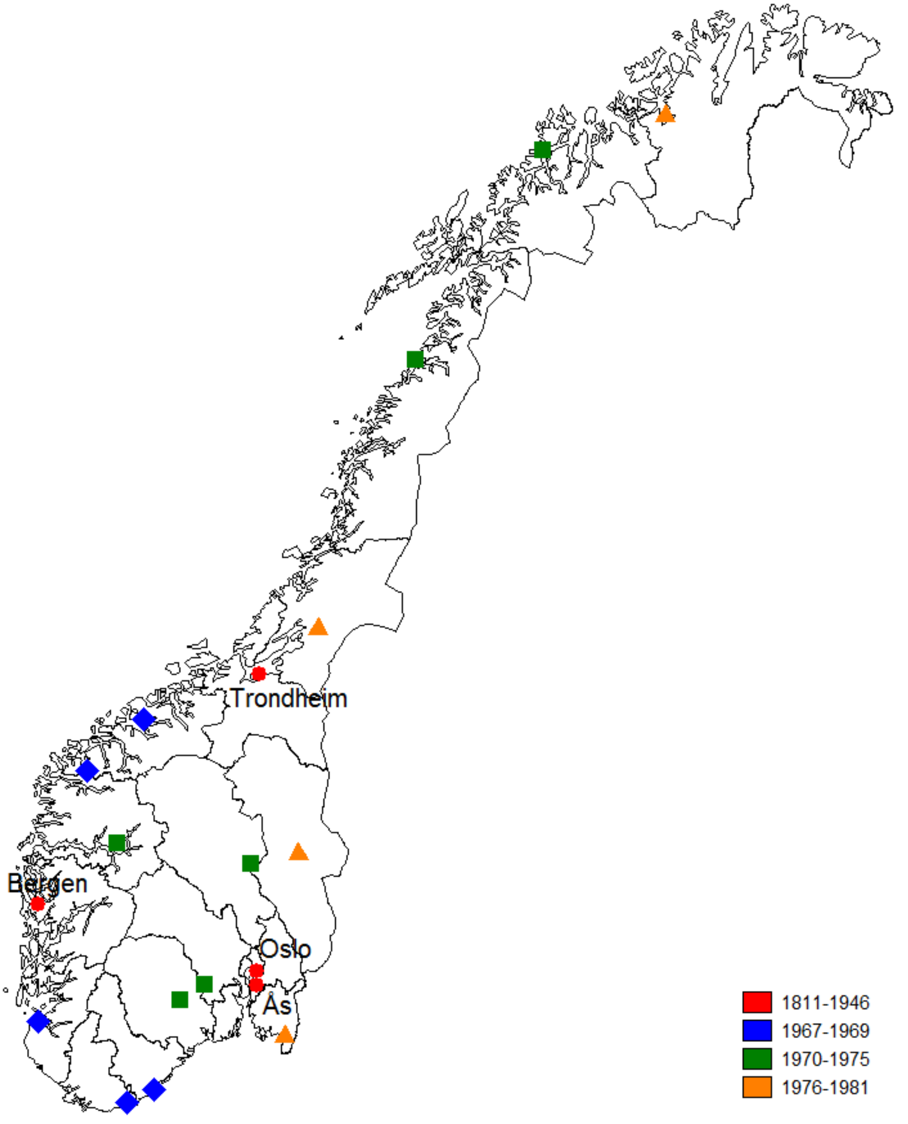

Note: This figure shows the geographic location of colleges across the country. We distinguish the colleges by their establishment year. The four universities established before the reform are labeled in red circles: they are located in Olso, Bergen, Trondheim and Ås. The remaining colleges shown on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

FIGURE A3. Relative Size of Full-time Students Enrollment, by Years since Establishment

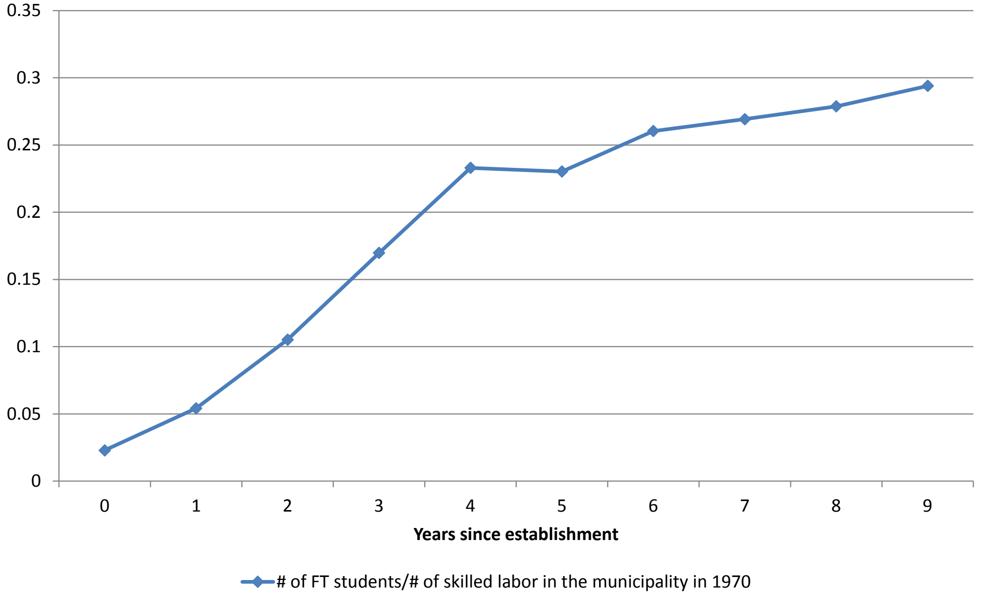

Note: This figure shows the total number of full-time students enrolled in the seven new regional colleges (opened between 1969 and 1971) in each year after establishment as a percentage of the total number of college-educated labor in the municipalities in 1970. Year 0 is normalized as the establishment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

FIGURE A4. Relative Size of R&D Activities, by Year

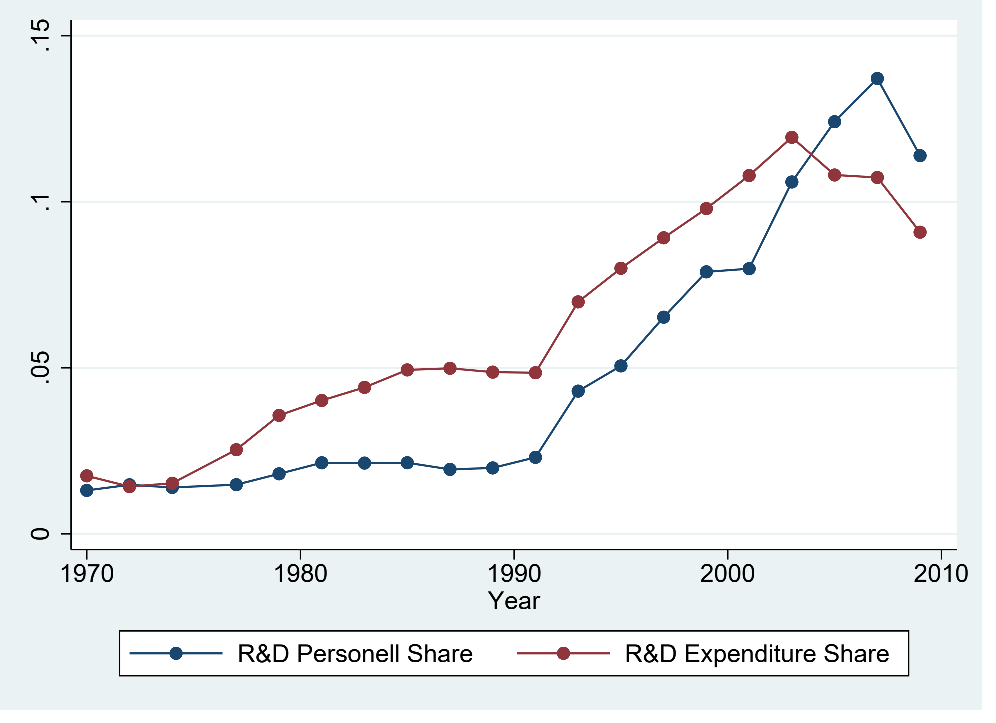

Note: This figure reports the shares of R&D personnel and R&D expenditures in the regional colleges relative to the totals in the higher education sector in Norway. Data source: Archival data provided by The Nordic Institute for Studies in Innovation, Research and Education.

FIGURE A5. The Effects of the Reform on Absolute Wages and Skill Compositions

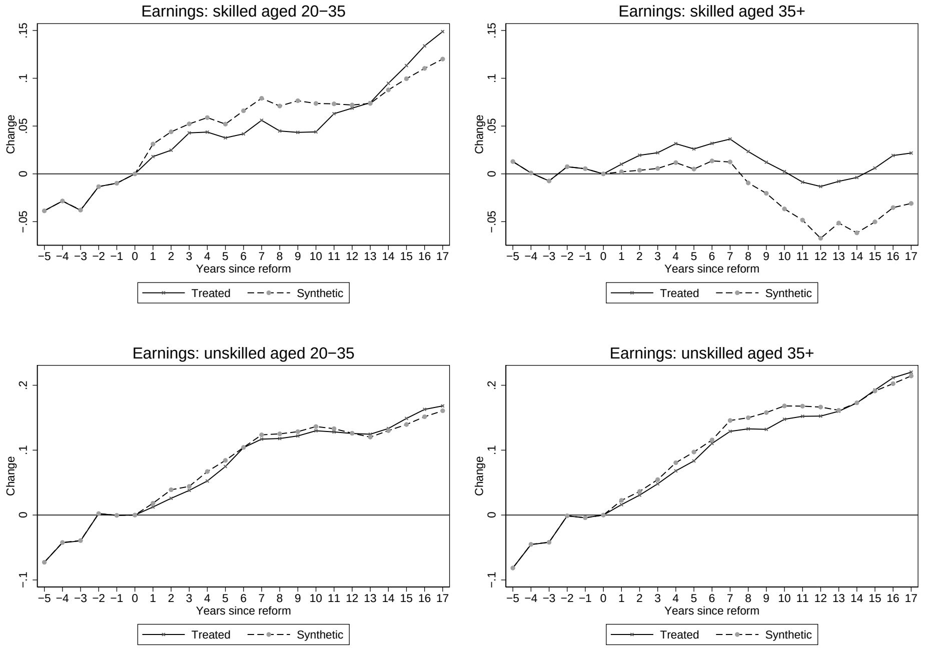

Note: This figure presents the synthetic control estimates on skill composition and absolute wages of the workforce. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome in a given period, relative to the levels in the year of the reform (where 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 A6. The Effects of the Reform on Relative Wages and Skill Compositions: Matching 2-year Pre-reform

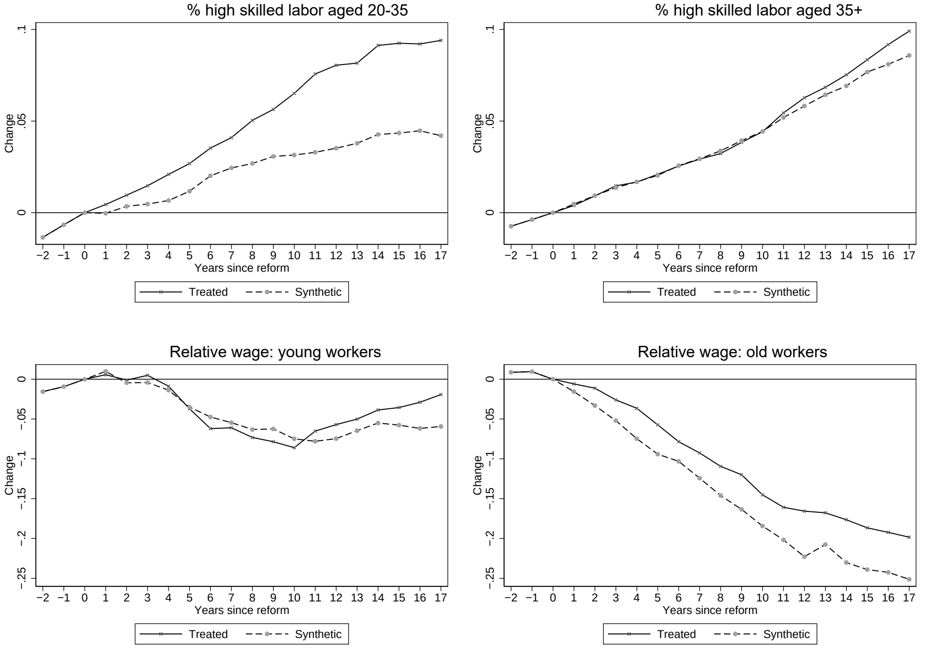

Note: This figure presents the synthetic control estimates on skill composition and relative wages of the workforce, where we match the outcomes for 1 and 2 years prior to the reform. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome in a 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 A7. The Effects of the Reform on Absolute Wages and Skill Compositions: Placebo Tests

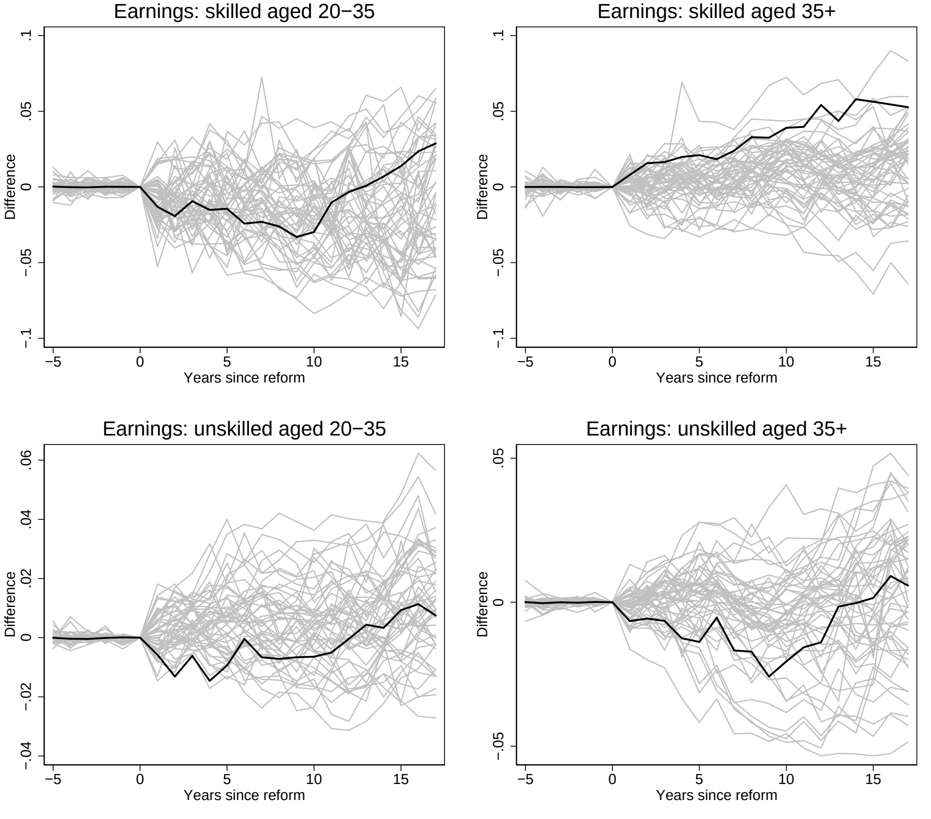

Note: This figure shows the results for the placebo test. The gray lines represent the estimated treatment effect from each placebo of the permutation test. The thick solid line denotes the treatment effect estimated using the actual treated municipalities in the data. The implied p-values of the actual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

FIGURE A8. The Effects of the Reform on Absolute Wages of Young Workers: STEM vs Non-STEM Colleges

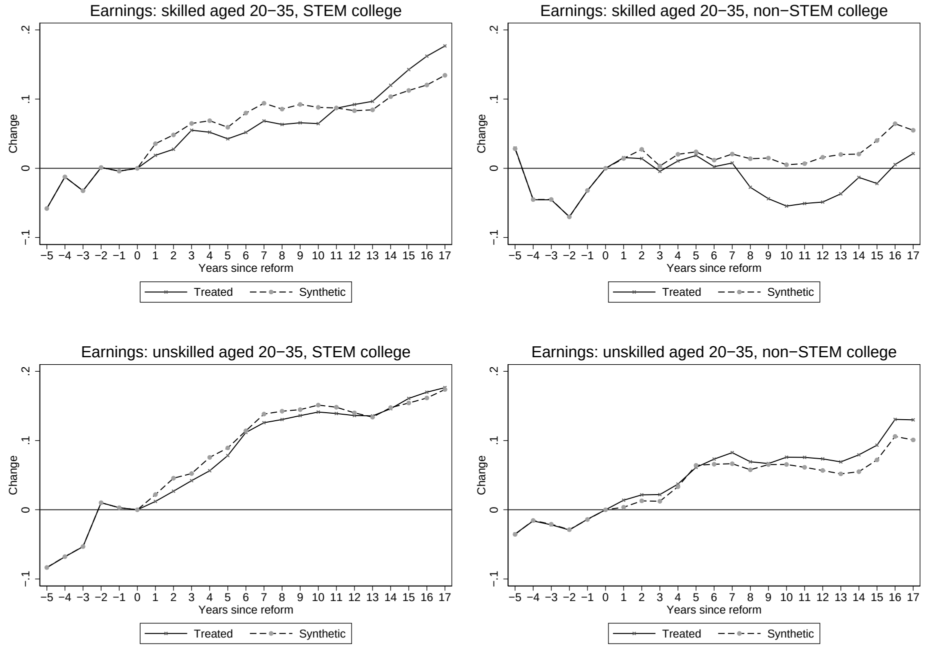

Note: This figure presents the synthetic control estimates on skill composition and wages of young workers, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome variable in a given period, relative 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 A9. The Effects of the Reform on Absolute Wages of Old Workers: STEM vs Non-STEM Colleges

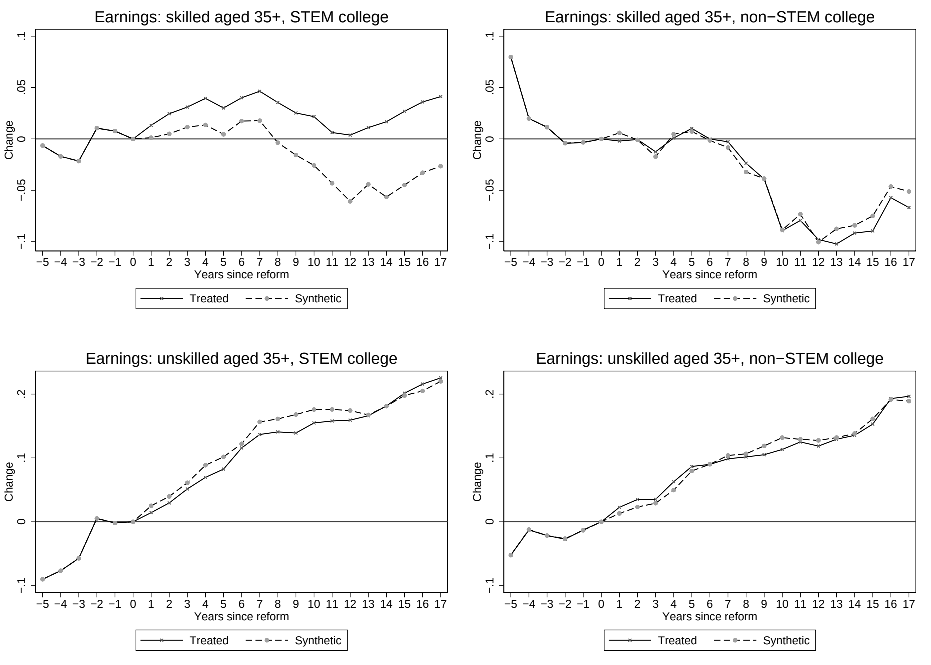

Note: This figure presents the synthetic control estimates on skill composition and wages of old workers, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome variable in a given period, relative 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 A10. Prediction for the Skill Wage Premium

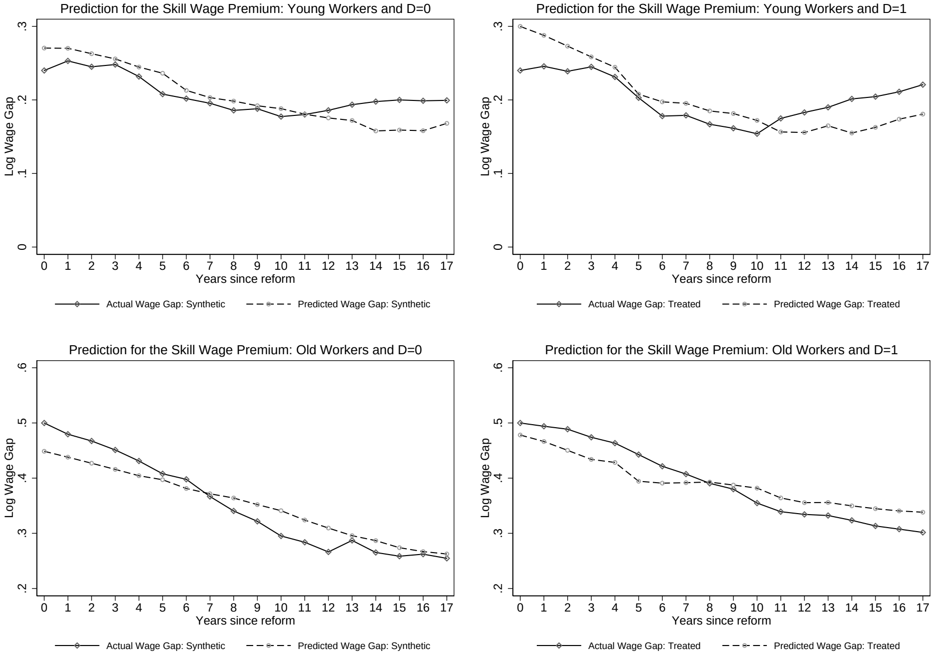

Note: This figure presents the fitness of the relative demand model (equation E.5) to the skill premiums of the treated group and synthetic control group estimated by the synthetic control analysis. See Section 3.3 for details.

FIGURE A11. Estimated Exogenous Technical Change and Total Technical Change, National Data 1967-1990

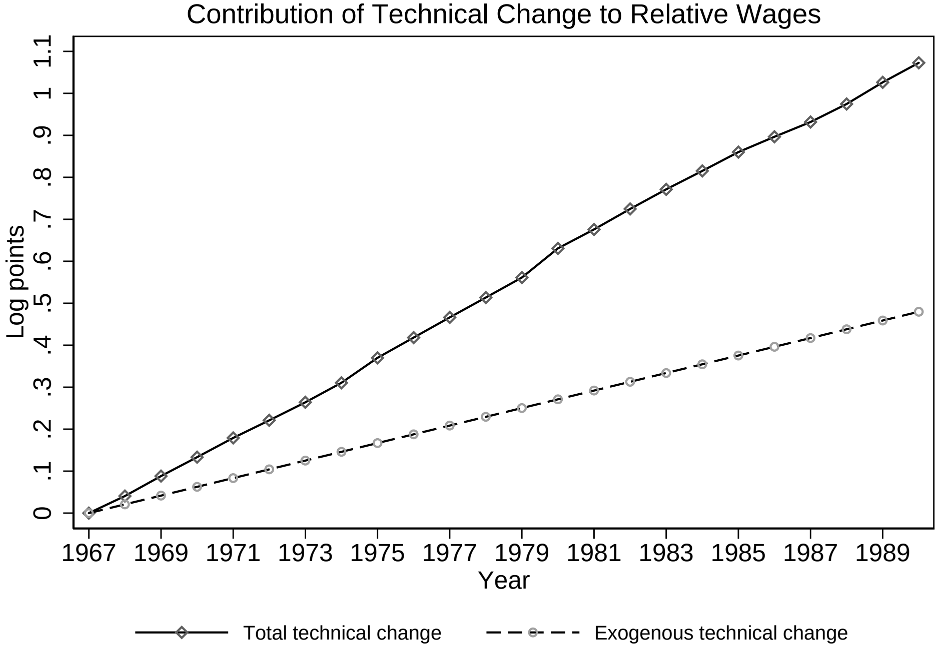

FIGURE A12. Total number of patents per year by regional colleges and universities, 1995-2010

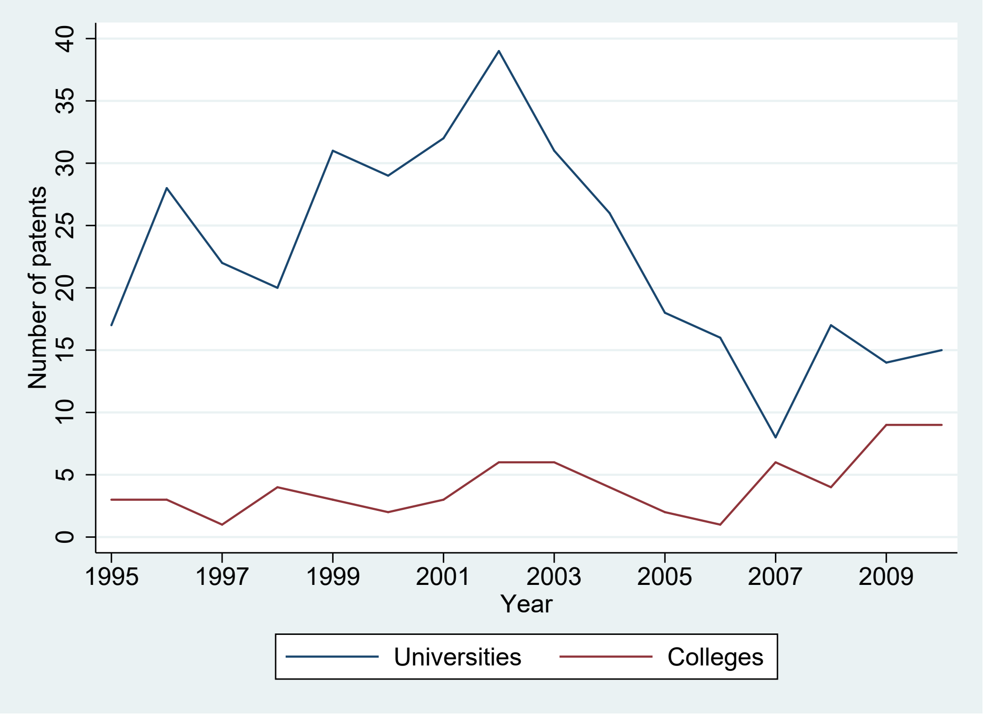

Data Source: Hvide and Jones (2018).

FIGURE A13. The Effects of the Reform on Inward Mobility

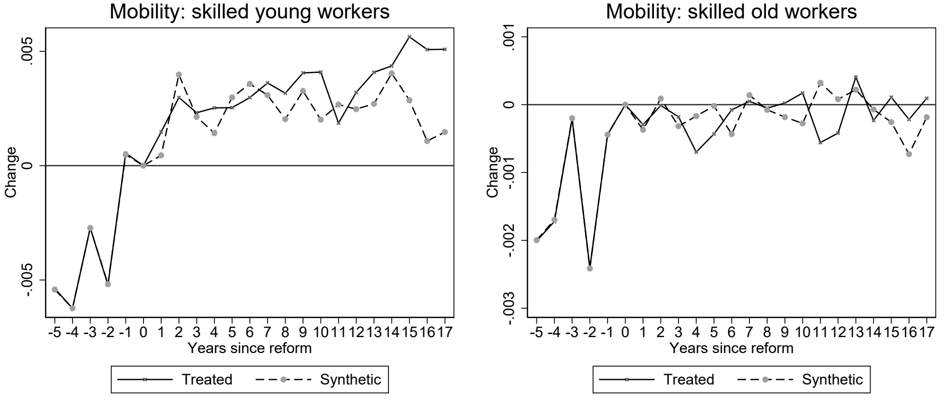

Note: This graph presents the synthetic control estimates on the share of skilled workers (by age group) moving into the municipality, relative to the population size of the receiving municipality. Each data point represents the mean outcome in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). The year of the reform is normalized to period zero.

FIGURE A14. The Effects of the Reform on Sector Compositions

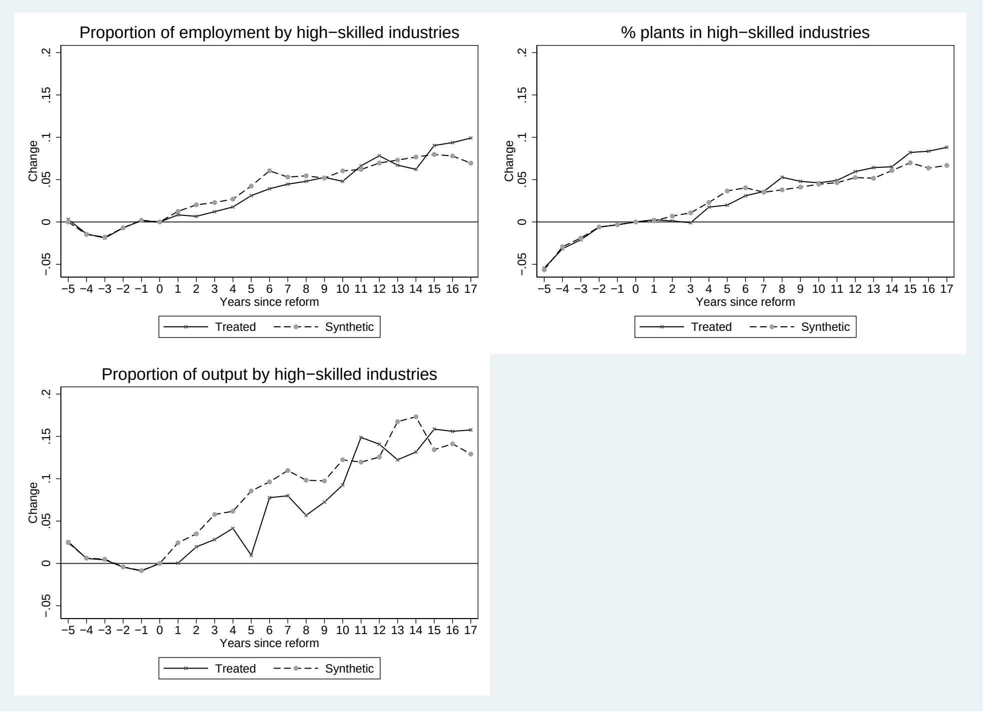

Note: This graph presents the synthetic control estimates on the share of output, employment and number of plants number from skill-intensive industries. Each data point represents the mean outcome in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). The year of the reform is normalized to period zero.

FIGURE A15. The Effects of the Reform on Value-added Output per Worker

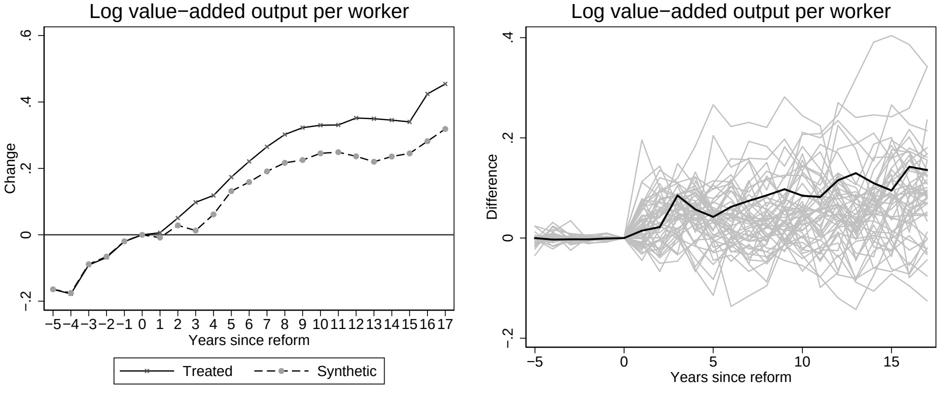

Note: The graph to the left presents the synthetic control estimates on log value-added output per worker. Each data point represents the mean outcome in a given period, relative to the levels in the year of the reform (where the outcome in the reform year is normalized to zero). The graph to the right 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 A16. The Effects of the Reform on Value-added Output per Worker: STEM vs Non-STEM Colleges

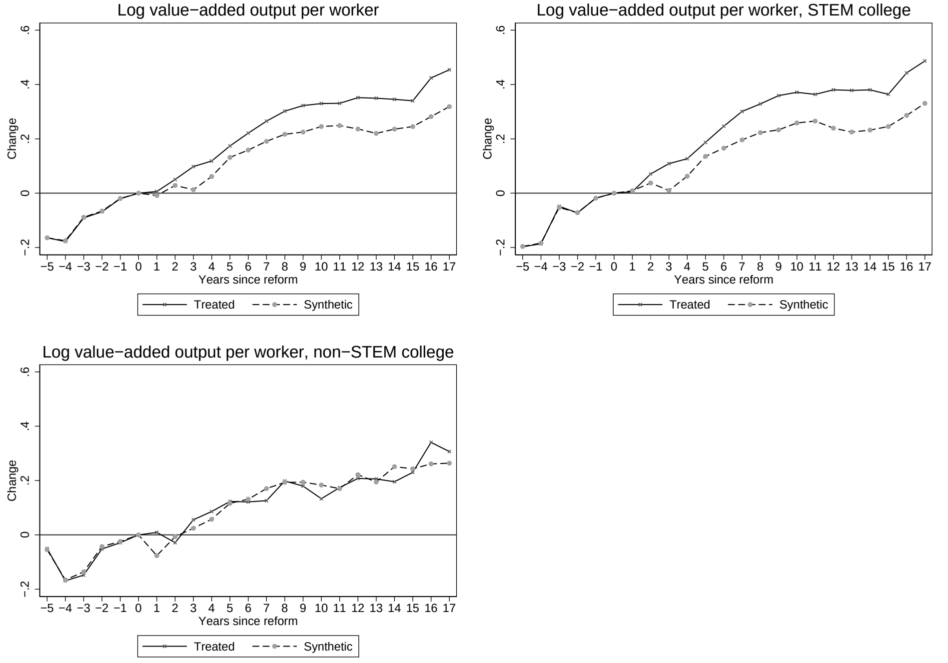

Note: This figure presents the synthetic control estimates on log value-added output per worker, separately for STEM colleges and non-STEM colleges. On each graph, the year of the reform is normalized to period zero. Each data point represents the mean outcome in a given period, relative to the levels 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	skilled worker %		Relative wages		Absolute wages			
post reform	Young	Old	Young	Old	S, young	S, old	U, young	U, old
1	.1	.54	.34	.32	.56	.28	.18	.16
2	.04	.3	.42	.1	.76	.06	0	.18
3	.04	.2	.46	.1	.46	.12	0	.18
4	.1	.36	.52	.06	.68	.08	.02	.14
5	.04	.28	.4	.1	.52	.08	.06	.1
6	.08	.38	.8	.18	.76	.14	.32	.24
7	.04	.38	.64	.18	.64	.12	.14	.16
8	.02	.38	.76	.14	.72	.08	.22	.22
9	.04	.32	.68	.12	.68	.06	.26	.12
10	.02	.28	.62	.14	.7	.04	.2	.16
11	.02	.14	.28	.16	.32	.06	.28	.22
12	0	.1	.3	.12	.4	.02	.34	.3
13	0	.14	.18	.12	.22	.06	.56	.46
14	0	.12	.18	.14	.24	0	.42	.42
15	0	.12	.14	.12	.2	.06	.54	.54
16	0	.08	.1	.16	.18	.04	.52	.52
17	0	.06	.24	.18	.22	.04	.36	.48
1–6	.04	.3	.54	.14	.62	.04	.06	.12
7–12	.02	.22	.5	.12	.6	.02	.24	.18
13+	0	.1	.14	.12	.2	.02	.44	.48

Note: This table reports the p-values implied by the permutation test of the synthetic control estimator. See Appendix Section B for details. The p-value in each period τ is calculated by $\frac{\sum_{l=1}^{K} \mathbf{1}(\widehat{\alpha}_\tau^l > \widehat{\alpha}_\tau^0)}{K}$, where K is the number of placebos (we use $K = 50$); $\widehat{\alpha}_\tau^l$ is the treatment effect from each placebo; and $\widehat{\alpha}_\tau^0$ is the estimated treatment effect (reported in Figure 2). For unskilled wages, we report the p-values using $\frac{\sum_{l=1}^{K} \mathbf{1}(\widehat{\alpha}_\tau^l < \widehat{\alpha}_\tau^0)}{K}$.

TABLE A2. 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	0.004**	-0.002	-0.009	-0.007**
	(0.002)	(0.012)	(0.013)	(0.003)
Years 3 to 4	0.008**	-0.006	-0.015	-0.010***
	(0.004)	(0.009)	(0.010)	(0.004)
Years 5 to 8	0.014***	-0.012	-0.020**	-0.008
	(0.004)	(0.009)	(0.010)	(0.006)
Years 9 to 12	0.030***	-0.020**	-0.027**	-0.007
	(0.006)	(0.008)	(0.012)	(0.009)
Years 13+	0.039***	0.002	-0.002	-0.004
	(0.008)	(0.014)	(0.020)	(0.011)
Pre-reform				
Years 2 to 1	-0.005***	-0.002	-0.002	-0.000
	(0.002)	(0.010)	(0.008)	(0.004)
Years 3 to 4	-0.006	-0.006	-0.008	-0.002
	(0.004)	(0.012)	(0.011)	(0.008)
Years 5 or before	0.010***	0.017	0.004	-0.012
	(0.004)	(0.014)	(0.009)	(0.013)
N	9127	9127	9127	9127

Note: This table reports estimates using equation (C.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 A3. 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**	0.008	0.004	-0.004
	(0.001)	(0.006)	(0.005)	(0.003)
Years 3 to 4	0.004***	0.011	0.005	-0.006
	(0.001)	(0.007)	(0.006)	(0.004)
Years 5 to 8	0.005***	0.017**	0.008	-0.008*
	(0.002)	(0.008)	(0.007)	(0.005)
Years 9 to 12	0.010***	0.022**	0.015*	-0.008
	(0.003)	(0.009)	(0.008)	(0.006)
Years 13+	0.021***	0.018	0.016	-0.002
	(0.004)	(0.012)	(0.013)	(0.011)
Pre-reform				
Years 2 to 1	-0.001	-0.004	-0.002	0.002
	(0.001)	(0.007)	(0.006)	(0.003)
Years 3 to 4	-0.009***	-0.005	-0.005	-0.000
	(0.003)	(0.012)	(0.009)	(0.010)
Years 5 or before	-0.002	0.017	0.007	-0.010
	(0.005)	(0.017)	(0.009)	(0.013)
N	9127	9119	9119	9127

Note: This table reports estimates using equation (C.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. Implied p-values from the Permutation Tests: Investment in Machinery and Equipment

# of years	Investment in		
post reform	machinery	machinery ex. Transport	machinery and facility
1	.44	.74	.46
2	.54	.5	.46
3	.18	.24	.66
4	.38	.5	.5
5	.4	.42	.14
6	.48	.66	.66
7	.76	.76	.56
8	.46	.62	.7
9	.48	.34	.36
10	.68	.72	.38
11	.28	.28	.28
12	.36	.38	.14
13	.08	.14	.26
14	.24	.24	.48
15	.14	.26	.24
16	.22	.14	.08
17	.1	.12	.12
1–6	.46	.6	.46
7–12	.5	.54	.22
13+	.06	.08	.14

Note: This table reports the p-values implied by the permutation test of the synthetic control estimator. See Appendix Section B for details. The p-value in each period τ is calculated by $\frac{\sum_{l=1}^{K} \mathbf{1}(\widehat{\alpha}_\tau^l > \widehat{\alpha}_\tau^0)}{K}$, where K is the number of placebos (we use $K = 50$); $\widehat{\alpha}_\tau^l$ is the treatment effect from each placebo; and $\widehat{\alpha}_\tau^0$ is the estimated treatment effect (reported in Figure 2).

TABLE A5. Local Labor Demand Estimation: First-step Estimates

	(1)	(2)	(3)
Age-group specific	-0.301***	-0.329***	-0.196***
relative supply	(0.032)	(0.033)	(0.042)
Year effects:			
Year 1	0.008	0.001	-0.007
	(0.015)	(0.012)	(0.004)
Year 2	0.016	0.006	-0.010**
	(0.015)	(0.012)	(0.004)
Year 3	0.022	0.015	-0.008*
	(0.015)	(0.012)	(0.004)
Year 4	0.027*	0.013	-0.015***
	(0.015)	(0.012)	(0.004)
Year 5	0.030*	0.017	-0.013***
	(0.015)	(0.012)	(0.004)
Year 6	0.013	0.009	-0.004
	(0.015)	(0.012)	(0.004)
Year 7	0.026	0.013	-0.013***
	(0.015)	(0.012)	(0.004)
Year 8	0.036**	0.021	-0.015***
	(0.015)	(0.012)	(0.004)
Year 9	0.040**	0.021	-0.019***
	(0.015)	(0.013)	(0.004)
Year 10	0.048***	0.031**	-0.017***
	(0.015)	(0.013)	(0.004)
Year 11	0.065***	0.049***	-0.016***
	(0.016)	(0.013)	(0.004)
Year 12	0.076***	0.063***	-0.013***
	(0.016)	(0.013)	(0.004)
Year 13	0.062***	0.058***	-0.004
	(0.016)	(0.013)	(0.004)
Year 14	0.076***	0.072***	-0.005
	(0.016)	(0.013)	(0.004)
Year 15	0.077***	0.076***	-0.001
	(0.016)	(0.013)	(0.004)
Year 16	0.078***	0.081***	0.003
	(0.016)	(0.013)	(0.004)
Year 17	0.091***	0.089***	-0.001
	(0.016)	(0.013)	(0.004)
N	36	36	36

Note: Column (1) of this table reports the estimates from equation (E.2) in the text. As independent variables, the regression includes a constant, time dummies and differences in the relative supply index within age and between treatment groups. Columns (2) and (3) reports the estimates from equation E.10 and E.11, respectively. As independent variables, the regressions include a constant, time dummies and differences in the supply index within age and between treatment groups.

TABLE A6. Estimates from the Relative Labor Demand Regression: Allowing σ_A to Differ by Skill Groups

	(1)	(2)	(3)	(4)
Aggr. Supply	-0.548	-0.545	-0.540*	-0.491
	(0.345)	(0.338)	(0.322)	(0.309)
Trend	0.006	0.006	0.006	0.005
	(0.010)	(0.010)	(0.010)	(0.009)
Trend×Post	0.008**	0.009**	0.010**	0.011***
	(0.004)	(0.004)	(0.004)	(0.004)
Treated	0.015	0.021	0.027*	0.029**
	(0.018)	(0.016)	(0.015)	(0.014)
Post	-0.024	-0.038*	-0.057***	-0.075***
	(0.021)	(0.020)	(0.021)	(0.022)
Older worker	0.141***	0.141***	0.141***	0.141***
	(0.007)	(0.007)	(0.007)	(0.007)
Constant	-2.215***	-2.218***	-2.224***	-2.274***
	(0.351)	(0.343)	(0.327)	(0.314)
N	72	72	72	72

Note: This table reports the estimates of equation (E.12). See Appendix Section E.1 for details. As independent variables, each regression includes a constant, a time trend (t), time trend interacted with the post dummy ($t \times P_{t(D)}$), a treatment group indicator (D), a post dummy ($P_{t(D)}$), an age group indicator (b_j), and the relative supply index. In column (1), we set M = 2. In columns (2) to (4), we set M = 3, 4, and 5, respectively.

TABLE A7. Estimates from the Relative Wage and Supply Regressions Using National Data

	(1)	(2)	(3)
Age-specific Supply	-0.439***		
	(0.039)		
Aggr. Supply		-0.757***	-1.457***
		(0.179)	(0.179)
Trend		0.021***	0.021***
		(0.007)	(0.007)
N	48	48	48

Note: See Appendix Section G for details.

TABLE A8. Estimates of σ_E and σ_A in the Literature

		Estimated elasticity of substitution	
Paper	Data and sample	σ_E	σ_A
Katz and Murphy (1992)	US: CPS data, 1963-1987	1.4	3.3
Card and Lemieux (2001)	US: Census and March CPS, 1959-1996 Canada: Census 1980-1995 UK: GHS, 1974-1996	All three countries: in the range of 2 to 2.5	All three countries: in the range of 4 to 6
Autor, Katz, and Kearney (2008)	US: CPS data, 1963-2005	1.57	3.55
D'Amuri, Ottaviano, and Peri (2010)	Germany: IAB employment sample, 1987-2001	2.9	3.3
Brücker and Jahn (2011)	Germany: IAB employment sample, 1975-2004	3	8.6
Ottaviano and Peri (2012)	US: Census and March CPS	2	5
Glitz and Wissmann (2017)	Germany: IAB employment sample, 1980-2008	1.6	8.2

Note: σ_E is the elasticity of substitution between skill groups, and σ_A is the elasticity of substitution between age groups. For papers using the US CPS data, skilled group is some college and unskilled group is high school equivalent. For papers using German data, skilled workers include those with tertiary degree and unskilled workers include those with lower secondary and vocational training aggregate. Aggregation of age groups differs across these studies, ranging from 2 to 8 age groups.

TABLE A9. Implied p-values from the Permutation Tests: Sectoral Compositions and Output

# of years post reform	High-skilled industries employment	plants	output	Value-added per worker
1	.58	.54	.84	.52
2	.58	.44	.56	.68
3	.56	.68	.64	.24
4	.54	.52	.62	.4
5	.56	.64	.9	.52
6	.58	.5	.58	.28
7	.54	.44	.54	.2
8	.5	.28	.66	.18
9	.44	.38	.52	.14
10	.54	.38	.54	.3
11	.42	.38	.18	.28
12	.44	.54	.24	.2
13	.56	.3	.64	.1
14	.62	.46	.5	.14
15	.44	.44	.22	.4
16	.4	.28	.18	.24
17	.38	.34	.26	.22
1–6	.6	.58	.68	.38
7–12	.5	.4	.42	.22
13+	.44	.4	.34	.18

Note: This table reports the p-values implied by the permutation test of the synthetic control estimator. See Appendix Section B for details. The p-value in each period τ is calculated by $\frac{\sum_{l=1}^{K} \mathbf{1}(\widehat{\alpha}_\tau^l > \widehat{\alpha}_\tau^0)}{K}$, where K is the number of placebos (we use $K = 50$); $\widehat{\alpha}_\tau^l$ is the treatment effect from each placebo; and $\widehat{\alpha}_\tau^0$ is the estimated treatment effect (reported in Figure 2).

TABLE A10. College Reform, College Education, and Average IQ Scores by Skill Groups

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

Note: Estimates using equation (D.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 A11. Production Function Estimates: First Stage of the Control Function

	(1)	(2)	(3)	(4)	(5)
		M=2	M=3	M=4	M=5
D_{ct}	0.001				
	(0.003)				
$D_{ct} \times t_D$	0.002***				
	(0.000)				
$t_D \times P_{t(D)}$		0.002***	0.002***	0.002***	0.002***
		(0.000)	(0.000)	(0.000)	(0.000)
$P_{t(D)}$		0.000	-0.000	-0.001	-0.001
		(0.004)	(0.004)	(0.004)	(0.004)
$D_{ct} \times (1 - P_{t(D)})$		0.003	0.004*	0.005**	0.005**
		(0.002)	(0.002)	(0.002)	(0.002)
Test of $D_{ct} \times (1 - P_{t(D)}) = 0$ (F-stat)		2.32	3.38	4.21	4.5
N	18439	18439	18439	18439	18439

Note: This table reports the first-stage of the control function estimates of production function. Variable D_{ct} is equivalent to the reform indicator. t_D is the number of years since the college opening (=0 in the years up to the college opening, or if there was no college opening in the municipality). $P_{t(D)}$ is an indicator function that takes value 1 if $t_D \geq M$, where M=2, 3, 4 and 5 in columns (2)-(5), respectively. The exclusion restriction (which is included only in the first stage but not in the production function) is $D_{ct} \times (1 - P_{t(D)})$. For instance, if M=2 (column 2), the excluded variable is an indicator function that takes value 1 if the observation is within the first two years of the reform. Standard errors are clustered at the municipality level and given in parentheses.

TABLE A12. Production Function Estimates: Robustness Checks

	(1)	(2)	(3)
	M=3	M=4	M=5
Output elasticity: skilled labor			
Average growth per year	0.007***	0.008***	0.009***
	(0.002)	(0.002)	(0.002)
Output elasticity: unskilled labor			
Average growth per year	-0.010***	-0.010***	-0.011***
	(0.003)	(0.003)	(0.004)

Note: This table reports the predicted differences in the annual growth of output elasticity (starting from M years after the reform, by skill group) between the treated group and the control group (holding labor inputs at the mean). The estimating production function takes the form of Cobb-Douglas production function, under different assumptions of M (equation (I.2) in Appendix Section I). Number of observations = 18441. Standard errors are clustered at municipality level and given in parentheses.

TABLE A13. Production Function Estimates: Exogenous Supply

Variable	(1)	(2)
$log(S)\times$0–1 years post-reform	0.008	-0.011
	(0.024)	(0.022)
$log(S)\times$post	0.010	-0.018
	(0.034)	(0.031)
$log(S)\times$post$\times$trend	0.007***	0.007**
	(0.002)	(0.003)
$log(U)\times$0–1 years post-reform	-0.034	-0.012
	(0.028)	(0.027)
$log(U)\times$post	0.003	0.019
	(0.035)	(0.037)
$log(U)\times$post$\times$trend	-0.009***	-0.008*
	(0.003)	(0.004)
0–1 years post-reform	0.375	0.262
	(0.257)	(0.256)
post	-0.046	0.022
	(0.334)	(0.330)
post$\times$trend	0.036	0.023
	(0.025)	(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.

TABLE A14. Numerical Exercise: The Impacts of Technical Change and Factor Supplies on Factor Shares

Year	SH'_{Sct}	SH'^{*}_{Sct} $\rho = 0.5$	SH'^{*}_{Sct} $\rho = 0.9$
0	0.441	0.441	0.441
1	0.441	0.444	0.447
2	0.440	0.446	0.450
3	0.449	0.448	0.454
4	0.459	0.450	0.456
5	0.468	0.452	0.460
6	0.477	0.452	0.460
7	0.486	0.453	0.462
8	0.495	0.459	0.473
9	0.505	0.462	0.478
10	0.514	0.468	0.490
11	0.523	0.476	0.504
12	0.533	0.479	0.509
13	0.542	0.479	0.509
14	0.552	0.483	0.516
15	0.561	0.485	0.519

Note: Column (1) reports actual change in factor shares as implied by the Cobb-Douglas estimates. Columns (2) and (3) shows, for a given ρ, the implied change in the share in the absence of technical change.

References for Appendix

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. (2007): "Equilibrium bias of technology," *Econometrica*, 75(5), 1371–1409.

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.

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.

BRÜCKER, H., AND E. J. JAHN (2011): "Migration and wage-setting: reassessing the labor market effects of migration," *Scandinavian Journal of Economics*, 113(2), 286–317.

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.

D'AMURI, F., G. I. OTTAVIANO, AND G. PERI (2010): "The labor market impact of immigration in Western Germany in the 1990s," *European Economic Review*, 54(4), 550–570.

GLITZ, A., AND D. WISSMANN (2017): "Skill Premiums and the Supply of Young Workers in Germany," .

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.

HVIDE, H. K., AND B. F. JONES (2018): "University Innovation and the Professor's Privilege," *American Economic Review*, 108(7), 1860–98.

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.

KMENTA, J. (1967): "On Estimation of the CES Production Function," *International Economic Review*, 8(2), 180–189.

LEVINSOHN, J., AND A. PETRIN (2003): “Estimating production functions using inputs to control for unobservables,” *The Review of Economic Studies*, 70(2), 317–341.

OTTAVIANO, G. I., AND G. PERI (2012): “Rethinking the effect of immigration on wages,” *Journal of the European economic association*, 10(1), 152–197.

STATISTICS NORWAY, N. (1994): “Historisk statistikk 1994,” *Historical statistics*.

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): “Evneprøving av utskrivingspliktige i Norge 1950-53,” Arbeidsrapport nr. 26, INAS (in Norwegian).