Essays on Technology and Work

by

Joonas Tuhkuri

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

Author . . . .

Department of Economics
May 12, 2022

Certified by . . . .

Daron Acemoglu
Institute Professor
Thesis Supervisor

Certified by . . . .

David Autor
Ford Professor of Economics
Thesis Supervisor

Certified by . . . .

Simon Jäger
Assistant Professor of Economics
Thesis Supervisor

Accepted by . . . .

Abhijit Banerjee
Ford International Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Essays on Technology and Work

by

Joonas Tuhkuri

Submitted to the Department of Economics
on May 12, 2022, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis consists of four papers on technology, work, skills, and personality using novel large-scale data and methods. The first paper (Chapter 1, with Johannes Hirvonen and Aapo Stenhammar) presents novel evidence on the effects of advanced technologies on employment, skill demand, and firm performance. The main finding is that advanced technologies led to increases in employment and no change in skill composition. Our main research design focuses on a technology subsidy program in Finland that induced sharp increases in technology investment in manufacturing firms. Our data directly measure multiple technologies and skills and track firms and workers over time. We demonstrate novel text analysis and machine learning methods to perform matching and to measure specific technological changes. To understand our findings, we outline a theoretical framework that contrasts two types of technological change: process versus product. We document that the firms used new technologies to produce new types of output rather than replace workers with technologies within the same type of production. The results contrast with the ideas that technologies necessarily replace workers or are skill biased.

The second paper (Chapter 2, with Ramin Izadi) investigates which personality traits and skills help workers to deal with a changing environment. Labor markets are in constant change. This paper documents how responses to labor-market shocks vary by individuals' psychological traits. We construct measures of cognitive ability, extraversion, and conscientiousness using standardized personality and cognitive tests administered during military service to 79% of Finnish men born 1962–1979. We analyze establishment closures and mass layoffs between 1995–2010 and document heterogeneous responses to the shock. Extraversion is the strongest predictor of adaptation: the negative effect of a mass layoff on earnings is 20% smaller for those with one standard deviation higher scores of extraversion. Conscientiousness appears to have no differential impact conditional on other traits. Cognitive ability and education predict a significantly smaller initial drop in earnings but have no long-term advantage. Our findings appear to be driven directly by smaller disemployment effects: extraverted and high cognitive-ability individuals find re-employment faster in a similar occupation and industry they worked in before. Extraversion's adaptive value is robust to controlling for pre-shock education, occupation, and industry, which rules out selection into different careers as the driving mechanism. Extraverts are slightly more likely to retain employment in their current establishment during a mass layoff event, but the retention effect is not large enough to explain the smaller earnings drop.

The third paper (Chapter 3, with Ramin Izadi) explores how different dimensions of personality predict school vs. labor-market performance, and how the value of these traits changed over time. We answer these questions using data that includes multidimensional personality and cognitive test scores from mandatory military conscription for approximately 80% of Finnish men. We document that some dimensions of noncognitive skills are productive at school, and some dimensions are counterproductive at school but still valued in the labor market. Action-oriented traits (activity,

sociability, and masculinity) predict low school performance but high labor market performance. School-oriented traits, such as dutifulness, deliberation, and achievement striving, predict high school performance but are not independently valued in the labor market after controlling for school achievement. We further document that the labor-market premium to action-oriented personality traits has rapidly increased over the past two decades. To interpret the empirical results, we outline a model of multidimensional skill specialization. The model and evidence highlight two paths to labor-market success: one through school-oriented traits and formal skills, and one through action-oriented traits and informal skills.

The fourth paper (Chapter 4) analyzes the impact of manufacturing decline on children. To do so, it considers local employment structure—characterizing lost manufacturing jobs and left-behind places—high-school dropout rates, and college access in the US over 1990–2010. To establish a basis for causal inference, the paper uses variations in trade exposure from China, following its entry to the WTO, as an instrument for manufacturing decline in the US. While the literature on job loss has emphasized negative effects on children, the main conclusion of this research is that the rapid US manufacturing decline decreased high-school dropout rates and possibly increased college access. The magnitudes of the estimates suggest that for every 3-percentage-point decline in manufacturing as a share of total employment, the high-school dropout rate declined by 1 percentage point. The effects are largest in the areas with high racial and socioeconomic segregation and in those with larger African American populations. The results are consistent with the idea that the manufacturing decline increased returns and decreased opportunity costs of education, and with sociological accounts linking the working-class environment and children's education.

JEL Codes: J00, J24, O33.

Thesis Supervisor: Daron Acemoglu
Title: Institute Professor

Thesis Supervisor: David Autor
Title: Ford Professor of Economics

Thesis Supervisor: Simon Jäger
Title: Assistant Professor of Economics

Acknowledgments

I'd like to give thanks to:

Advisors: David Autor for advice, support, kindness, humor, and being miraculously available. Daron Acemoglu for always keeping his eyes on the big picture and showing how to do it. Simon Jäger for jumping in, sensibility, and guidance in the final steps.

Co-authors: Johannes Hirvonen, Ramin Izadi, and Aapo Stenhammar.

Funders: MIT, the Jerry Hausman, Stanley and Rhoda Fischer, Castle Krob International, and Center for International Studies Fellowships, the Harvard CES Dissertation Completion Fellowship, Etla Economic Research, Research Division on Business Subsidies (Ministry of Economic Affairs and Employment), George and Obie Shultz Fund, Kone Foundation, Emil Aaltonen Foundation, Foundation for Economic Education, Labour Foundation, Kordelin Foundation, Otto Malm Foundation, Oskar Öflund Foundation, OP-Group Foundation, and Yrjö Jahnsson Foundation.

Data and Fieldwork: Statistics Finland, the Ministry of Economic Affairs and Employment, particularly Samuli Neuvonen, Satu Nurmi, and Valtteri Raudaskoski. Ahmed Hewidy, Vertti Neuvonen, and Veeti Vallin at Etla for expert research assistance. The ELY Centers, Industrial Union, firm representatives, and workers who talked to me.

Institutions: MIT, Etla Economic Research and Antti Kauhanen (for making this project possible), Harvard Center for European Studies, Aalto University, the US Embassy, and the Fulbright Program. Workshops at MIT, Harvard, Etla, Helsinki GSE, BI, Boston U, Columbia GSB, Humboldt, LMU, Stockholm, Uppsala, and Wisconsin. Development Tea and Labor Coffee. The Economist, Wired, Helsingin Sanomat, YLE, and Kauppalehti.

Feedback: Alberto Abadie, Joshua Angrist, Anu-Hanna Anttila, David Atkin, Aleksi Avela, Manuel Bagues, Pedro Bessone, Suzanne Berger, Riku Buri, Jonathan Cohen, Mert Demirer, Ciprian Domnisoru, Dave Donaldson, Benjamin Faber, Robert French, Chishio Furukawa, Flavio Hafner, John Horton, Markus Jokela, Martti Kaila, Tuomas Kosonen, Felix König, Matti Mitrunen, Jonas Mueller-Gastell, Benjamin Olken, Teemu Pekkarinen, Tuomas Pekkarinen, Michael Piore, Drazen Prelec, Charlie Rafkin, Gautam Rao, Daniel Rock, Tobias Salz, Matti Sarvimäki, Karthik Sastry, Frank Schilbach, Mikko Silliman, Jeffrey Smith, Kosti Takala, Marko Terviö, Otto Toivanen, Roope Uusitalo, John Van Reenen, Briitta van Staalduinen, Sean Wang, Michael Wong, Samuel Young, and particularly Bengt Holmström (for encouragement, connections, and conversations).

Before: Philip Oreopoulos and Petri Rouvinen for helping me get to MIT.

Family: Jukka Tuhkuri, Maaria Linko, Anna Tuhkuri Matvejeff, and Ville Matvejeff, Ritva and Matti Linko, and Irja Tuhkuri. And Maija Greis for supporting this project and for far more.

Contents

Chapter 1

New Evidence on the Effect of Technology on Employment and Skill Demand

With Johannes Hirvonen and Aapo Stenhammar

1.1 Introduction

A central question in the debate on the future of work is: What are the effects of advanced technologies on employment and skill demand? Two ideas often dominate the conversation. The first is that technologies replace workers (the Luddites; Keynes 1931; Brynjolfsson and McAfee 2014). The second is that technologies increase the demand for skills and can increase inequality—this is called the skill-biased technological change hypothesis (Griliches 1969; Welch 1970; Tinbergen 1975). Current research suggests that advanced technologies such as robots and ICT have been skill biased (Katz and Murphy 1992; Krusell et al. 2000; Autor et al. 2003; Acemoglu and Autor 2011; Akerman et al. 2015; Acemoglu and Restrepo 2020). But the evidence is limited because both measuring and identifying the effects of technologies are difficult.

This paper presents novel evidence on the effects of advanced technologies on employment, skill demand, and firm performance using new large-scale data and quasi-experimental designs. The context is manufacturing firms in Finland, 1994–2018. We focus on new production technologies, such as robots and computer numerical control (CNC) machines. Our novel data directly measure technologies, employment, and skills and track firms and workers over time. The main research design focuses on a technology subsidy program that induced sharp increases in technology investment in specific firms. The program provides direct funding for technology investment and is part of the European Structural and Investment Funds—one of the world's largest industrial policy programs. Our design compares close winners and losers of the technology subsidies using an event-study approach. We use novel text analysis methods on the application text data to compare close winners and losers (meaning that the firms had similar evaluation reports) and measure specific technological changes (Roberts et al. 2020). We complement our quantitative analysis with fieldwork:

observing factories and interviewing CEOs, managers, workers, and subsidy administrators.

The first part of the paper reports results in sharp contrast with the ideas that technologies necessarily reduce employment or are skill biased. Technology investments induced by the subsidy program led to a 23% increase in employment, on average. But there were no differential changes in typical measures of skill bias: share of highly educated workers, average years of education, or production workers' share of employment. Zooming in to more detailed measures of skill composition—education and occupation groups, cognitive performance, and personality—we find generally zero effects. Several observations support the validity of our findings. The subsidy program induced a strong first stage: the firms showed a sharp rise in investments in technologies after winning technology subsidies. The firms had similar pre-trends in investment, employment, and skill composition before applying. Our results are robust to controlling for the evaluation texts of the subsidy applications using text analysis and other controls, including industry, firm size, and region trends. The results also hold when using alternative designs: a comparison to a matched non-applicant control group, a separate regression discontinuity (RD) design based on changes in the criteria defining a priority for small firms, and an event-study design without the subsidy program (Bessen et al. 2020). Our fieldwork supports these findings on the factory floor.

The second part of the paper explains the result that technologies did not replace workers or increase skill demand. To understand the findings, we outline a theoretical framework that contrasts two types of technological change: *process* versus *product*. The framework builds on Dixit and Stiglitz (1977) and Melitz (2003); we apply the ideas to a new context. Process refers to a productivity increase within an output variety, whereas product refers to expanding to new varieties. These two views predict different effects and can be tested empirically. The distinction is whether firms use new technologies to do the same thing at lower costs or do new things. The model clarifies that technologies may not necessarily be about changing the production process to replace workers or increase the demand for skills but creating new types of output. For example, automation is a process change, while the innovation of new goods is a product change (Klette and Kortum 2004; Acemoglu and Restrepo 2018).[1]

Based on the theoretical interpretation, we provide novel evidence documenting that the firms used technologies to create new products and services, not replace workers. Direct evidence shows that technology adoption led to more revenue, new products, and export growth. Text data from the subsidy program show that 91% of the firms described new products, response to changing demand, and other similar reasons for their technology investment. For example, the piston manufacturer included in the fieldwork invested in a new CNC machine and a robot to manufacture new, more effective pistons. Survey data from the EU's Community Innovation Survey (CIS) corroborate our observations: typical reasons for firms' process and product innovations are access to new markets, expanding product selection, and better quality—not typically to replace workers. We show the

[1]The concepts of process and product refer to the *uses* of technologies rather than physical *types* of technologies. Process, which is the idea that technological change lowers production costs, embeds the standard versions of labor replacement and skill bias. Conversely, product, which is the idea that technological change creates new output varieties, is present in standard growth models (Romer 1990; Grossman and Helpman 1991; Aghion and Howitt 1992) and in the management literature (Utterback and Abernathy 1975; Porter 1985).

results also hold without the subsidy program, indicating that our results are more general.

To understand when and why to expect process versus product changes, we contrast two types of manufacturing: *mass production* (Taylor, 1911; Ford, 1922) versus *flexible specialization* (Piore and Sabel 1984; Milgrom and Roberts 1990). Mass production combines standardized products, high volumes, and process advances. Flexible specialization combines specialized products, low volumes, and product advances. While the two ideas—labor replacement and skill bias—are widely accepted and used in the literature, research also recognizes that not all technological changes are labor replacing or skill biased. Most importantly, Piore and Sabel (1984) argue that a different set of technology–labor relations emerge in flexible manufacturing, most visible in technologically advanced small- and medium-sized enterprises that produce specialized products in small volumes to a changing market. In that context, and ours, the scope for specialization, low production volumes, and need for adaptation make it less profitable for firms to commit to the long-production runs of mass production and the fixed costs of process advances.[2] But our findings may not apply to non-specialized commodities, such as cement or steel, or high-volume assembly, where costs are critical. At the same time, the literature documents that manufacturing has widely evolved from mass production to flexible specialization (Dertouzos et al. 1989; Berger 2013).[3]

Two descriptive facts help position our findings into a broader context. First, the backdrop of our study is that the overall direction of manufacturing, including our treatment and control groups, is toward greater skill demand, seen in, for example, the rising share of educated workers. Because the skill trends are consistent with the rest of the world (Acemoglu and Autor 2011), we could have expected to find that new technologies were driving them at the firm level—but we did not. Our findings point to explanations for these skill trends other than the direct effects of adopting new technologies. Second, a critical aspect is that technology adopters are different from non-adopters. Growing firms typically invest in technologies, with and without subsidies. Our main design contrasts growing firms that plan to adopt new technologies. One firm gets the subsidy, the other does not, and that induces differences in technology adoption. This has two implications: 1) Our estimates capture the local average treatment effect (LATE) for firms close to investing in technologies. 2) Pre-screened but non-winning applicants provide a better control group than generic non-applicant firms because they have expressed an interest in technology adoption.

How broadly do the results apply? Our evidence is from Finland, where we can quantify the effects with high-quality data and research design. But the input we received from managers working in different contexts was that our observations apply more broadly in industrial manufacturing. There are still limitations. Our results do not directly apply to non-physical technological advances such as digitization or the internet, management practices such as lean manufacturing, R&D, technological advances in offices, historical eras, or the future. Our results and explanation focus on a firm-level mechanism. We do not exclude that micro-level technology could lead to macro-level skill

[2] Klette and Kortum (2004) and Akcigit and Kerr (2018) also relate the type of firm and innovation.

[3] Early research noted these changes first in Northern Italy, Germany, and Japan (Piore and Sabel 1984). Currently, the majority of Northern European manufacturing could be characterized as flexible specialization. For example, 90% of manufacturing employment in Finland is in non-commodity production under the Rauch (1999) classification. Bils and Klenow (2001) also document that US consumers have shifted away from standardized goods.

bias or labor replacement (Oberfield and Raval 2021). We also do not claim that work does not change: our qualitative evidence suggests it does, but that change does not imply labor replacement or skill bias by education, occupation, or cognitive performance.

Because our results challenge the two major ideas in the literature—that technologies replace labor or increase skill demand—it is critical to compare them to earlier research. We make two methodological contributions: We are the first to study the effects of technologies in manufacturing using a direct firm-level quasi-experiment, and our measurement is a major advance over earlier work because we directly measure the critical objects: technology, employment, and skills. Our results differ from the theoretical literature because it has focused more on process advances in mass production (Acemoglu and Restrepo 2018), while product advances are more common in our context. Our results are consistent with the non-quasi-experimental empirical studies that focus on similar technologies in manufacturing firms (Doms et al. 1997; Bartel et al. 2007; Aghion et al. 2020; Dixon et al. 2021; Koch et al. 2021) and qualitative evidence (Berger 2020). Complementary and simultaneous work by Curtis et al. (2021) documents that capital tax credits that favor capital investment raised labor demand in US manufacturing based on industry-level exposure. One interpretation is that their study detects similar local effects in the frontier sectors: their effects appear the largest in capital-intensive, skill-intensive, and robot-intensive subsectors of manufacturing. Potentially, capital subsidies made to frontier sectors are generally not applied to labor savings but rather market-share expansion among differentiated goods producers.[4]

Our analysis also contributes to the literature on industrial policy. We provide new estimates for one policy: a lump-sum transfer to increase technology adoption in manufacturing firms. The estimates help understand the broader question in growth and trade policy: What types of policies help firms grow? (Rodrik 2007). We find that the firms in our context use subsidies and technologies to achieve growth. To do so, they often scale up from idea to production. Our quantitative estimates suggest that 1 euro in technology subsidies led to 1.3 euros of technology investment. A typical EUR 100K subsidy led to 2.3 new jobs over the next 5 years. The cost per job was EUR 43K, close to the literature's average (Criscuolo et al. 2019).[5]

The paper proceeds in two parts. The first part presents the context, data, empirical strategies, and key results on employment, skill composition, and firm performance. The second part offers a theoretical interpretation based on process vs. product advances and then provides theory-motivated tests of that interpretation. Finally, we analyze robustness and conclude.

[4]Empirical studies also find different effects when focusing on 1) different types of technologies (especially digital technologies—the internet in Akerman et al. 2015 and ICT in Gaggl and Wright 2017), 2) replacement effects (Bessen et al. 2020), and 3) macro-level comparisons (Lewis 2011; Michaels et al. 2014; Acemoglu and Restrepo 2020).

[5]Recent research on industrial policy include Becker et al. (2010), Cerqua and Pellegrini (2014), Howell (2017), Criscuolo et al. (2019), Giorcelli (2019), Curtis et al. (2021), Howell et al. (2021), and Lane (2021). Technology subsidies and taxes are also actively debated (Acemoglu et al. 2020a; Costinot and Werning 2020; Guerreiro et al. 2021). We further review related research in Appendix A.10.

1.2 Context

We analyze the effects of advanced technologies in manufacturing firms in Finland, 1994–2018. Because we study technology investment with and without the subsidy program, we first outline the context common to all our analyses.

The technologies in our context are standard new production technologies in manufacturing: new CNC machines, robots, laser cutters, surface-treatment technologies, measurement devices, enterprise resource planning (ERP), computer-aided design (CAD) software, and similar technologies. The workers are primarily production workers (median 70%), for example, machinists, welders, and machine operators, typically with vocational training. The most represented industries are fabricated metal products and machinery. The firms are typically medium and small-sized (SMEs), but we also analyze large firms. Most firms are contract manufacturers that produce specialized intermediate goods in small batches, for example, pistons for engines, for large exporting firms. Figure 1-1 provides photographs of the typical technologies, workers, and firms in our sample.

Figure 1-2 documents that the overall direction of Finnish manufacturing is towards greater skill demands, seen in a rising share of educated labor and college income premium and a falling production-worker share. Finland's trends are consistent with the rest of the world (Acemoglu and Autor 2011), and the firm-level mechanisms we document might not be limited to Finland.

"Moore's Law for Pistons"

We conducted fieldwork to document the sample firms' technology adoption. The case of an industrial piston manufacturer clarifies our context.

The firm had invested in a new CNC machine, a robot arm, a measurement device, and new CAD software. When asked why they adopted the new technologies, the firm wanted to illustrate what they considered as the big picture of technological change in piston manufacturing: constant quality improvement. "With the old technologies, we couldn't make these pistons." Quality is essential for the piston manufacturer: pistons are only a fraction of an industrial engine's price, but if they break, it is expensive (see Kremer 1993 and Autor 2015 on the O-ring production function). Figure 1-3 shows the development of piston quality over the last 100 years. The firm called this the "Moore's law for pistons." The main effect of the new technology was that the firm could now produce new, larger, and more effective pistons. The firm stayed competitive and, as a result, has increased its revenue and employment.

The technology investment was associated with changes in production and work experience. Mainly those were "small, but important changes." For example, the new production design included a proprietary method of attaching the piston to the machining platform. The new production required some new skills: production workers needed to learn to use the robot and the CNC machine, and the R&D team had to learn to program with the new CAD software. The educational composition did not change as a result of the investment. But the educational composition in the firm has been increasing secularly over time.

The firm described operating in an environment where the market for each specific product is limited. They are de-facto monopolists (or oligopolists) in that market. They could not expand substantially within a product but could potentially expand by introducing a new product. All firms we studied explained essentially the same story, suggesting that the mechanisms could apply to other industrial and custom manufacturing firms.

1.3 Data

The first challenge in estimating the effects of technology on employment and skill demand is measurement. We directly measure the critical objects—technologies, work and skills, and firm performance—using novel high-quality data that track workers and firms over time.[6]

1.3.1 Technologies

We measure technologies using financial, text, customs, and survey data.

Financial Data[7] The primary source for measuring firms' technology investment is the Finnish Financial Statement Register. We measure firms' total investment and separately machinery and equipment and software. Statistics Finland collects the data directly, and the data cover all Finnish enterprises in almost all industries and our analysis years 1994–2018.

Text Data We develop a method to measure technologies using text data.[8] We measure overall technology investment, types of technologies, and uses of technologies directly at the firm level. The information on technologies' uses allows us to measure process vs. product advances.

The source for our text data is the ELY Center subsidy program, described in Section 1.4. The text data are unstructured and produced as a side product of the program. A technology subsidy application typically specifies the technology's *type* (e.g., a welding robot) and its *use* (e.g., weld longer seams). We focus on summary texts written by the program officers. The texts provide information on firms' actual plans because the technology plan is binding; the firms receive subsidies against verifiable costs. The full data contain 42,909 subsidy applications in different categories: technologies, exports, R&D, start-up, etc. Our method works in two steps:

Step 1: We code 21,210 randomly selected texts into categories based on pre-determined criteria, summarized in Table 1.1. We distinguish the type and use of technology because a firm can use the same technology for multiple purposes. Within technologies' uses, we code texts

[6] We provide details on data in Appendix A.5. For consistent measurement, we harmonize the Finnish occupation, industry, and geography classifications. The novel crosswalks are available at economics.mit.edu/grad/tuhkuri/data.

[7] We deflate all monetary values in this paper to 2017 euros using the Statistics Finland CPI.

[8] Many policy programs and firms' decisions leave a trail of text records. Using this method, researchers can use text to produce data retrospectively without new data collection and when data would not be available otherwise. The novel part of our research is to measure technologies directly within firms. Recent research uses text data to measure technological changes, especially patents, in other ways (Alexopoulos 2011; Atalay et al. 2020; Autor et al. 2021; Dechezlepretre et al. 2021; Howell et al. 2021; Kogan et al. 2020; Mann and Puttmann 2021; Webb 2020).

into applications intended to improve productivity within the same output variety (process) or produce new varieties (product). Within technologies' types, we code texts into automated vs. non-automated technologies (no active vs. an active user) and hardware vs. software (or both).

Step 2: We use machine learning to code the remaining 21,699 texts. We convert texts into a clean format, use the bag-of-words representation with TF-IDF weights, and support-vector machines (SVMs) for prediction. Figure A-46 presents features that best predict the technology category. Table A.45 provides summary information: our method achieves 95% accuracy in finding the technology applications from the pool of all applications. For the technology subcategories, we manually re-code all applications in the analysis sample to maximize precision.

Customs Data To measure the types of technologies, we also use customs data.[9] The data track technologies that firms import. Customs data record 621 different types of technologies in the 6-digit CN-classification system. We classify these technologies based on the physical type of machinery. The main distinction is between automated technologies vs. non-automated technologies. Automated technologies include, e.g., robots and CNC machines. Non-automated technologies include, e.g., non-automatic and hand-operated tools, hydraulic presses, and lifting equipment.

Survey Data To measure the uses of technologies, we also use survey data. The EU's Community Innovation Survey (CIS) provides firm-level information on the importance of different objectives for product and process innovations.

1.3.2 Work and Skills

We measure employment and wages from the registers maintained by Statistics Finland. The data allow us to track all individuals in Finland over time independently of their labor-market status. We link these data to multiple data sources on skills: education (level and type), school grades (9th grade GPA and high-school exit exam), and cognitive performance and personality (test scores from universal male conscription). We measure occupations from employment registers at the 3-digit level in the ISCO classification system. To measure the task content of occupations, we use the European Working Conditions Survey (EWCS) that provides information on the tasks workers perform in their jobs, collected through face-to-face interviews every five years. We construct occupation-level measures of task intensity for routine, manual, cognitive, and social tasks.

1.3.3 Firm Performance

We assemble a large set of data on firm performance, including revenue, productivity, profits, exports, products, prices, marketing, and patents. The data track all firms over time.

The firm-performance measures, revenue and profits, are obtained from Finnish Financial Statement Register. We use two variables to measure productivity: revenue per worker and total factor

[9] Recent research uses customs data to measure technology adoption; it is one of the few sources that track the types of technologies firms adopt (e.g., Acemoglu et al. 2020b; Acemoglu and Restrepo 2020, 2021).

productivity (TFP) estimated using the Cobb-Douglas production function.[10] We measure profits by the profit margin, defined as profits divided by revenue. We define the labor share as the wage bill divided by revenue. We winsorize firms' monetary values at the 5% level.

Exports are measured from Finnish Customs' Foreign Trade Statistics. We measure firms' products also from the Customs Register at the 6-digit CN classification. We focus on the number of products per firm and product turnover: introduced and discontinued products. We compute prices from the Customs Register and the Industrial Production Statistics, defining product-level prices as the product-level revenue divided by the number of units sold. Marketing expenditure data comes from the Financial Statement Register and patent data from Finnish Patent Database.

We measure firm subsidies from multiple registers. Two centralized systems (Yrtti 1 and 2) record the ELY Center subsidies. We gained access to these previously unstudied data that record the application process from submission to decision. We measure all other firm subsidies using the Statistics on Business Subsidies.

1.4 Research Design

The second challenge in estimating the effects of technology on employment and skill demand is identification. Our main research design is based on a technology subsidy program for manufacturing firms. Technology subsidies offer a valuable source of variation because they provide firms with a well-defined shock to the cost of technologies. We implement and validate an event-study design that compares close winning and losing firms of technology subsidies over time. The basis of the design is similar to Angrist (1998), Greenstone et al. (2010), and Kline et al. (2019).

A further novel aspect is that we use text data to create comparisons of close winners and losers. To do so, we use evaluation reports written by the program officers. We map these reports into propensity scores that reflect the likelihood of receiving a subsidy and control for the scores to compare close winners and losers. Roberts et al. (2020) discuss text matching.

We present two alternative designs in the Appendix: 1) a regression discontinuity (RD) design based on a change in the threshold that determines a priority for small firms in the program (to address internal validity), and 2) a spikes design based on the precise timing of technology adoption events without the program (to address external validity). These designs complement our overall argument, and we refer to them in the analysis.

1.4.1 The Subsidy Program

The Program The technology subsidy program is administrated in Finland by the Centers for Economic Development, Transport and the Environment (the ELY Centers).[11] These centers

[10] We obtain similar estimates using the Olley-Pakes and Levinsohn-Petrin methods (available upon request).

[11] There are 15 ELY Centers in our data. Until 2009 these centers were called TE Centers. Since 2014, four RR-ELY Centers have administrated all technology subsidies. ELY Centers are separate from Business Finland (previously TEKES), which provides funding for R&D.

promote regional business policy through various activities, including advisory, financing, and development services. Technology subsidies are part of a service called the Business Development Aid. The service provides funding for technology adoption, export promotion, R&D, and several smaller categories, such as starting a new company. It also supported firms during COVID-19. The service granted EUR 2 billion over our sample period 1994–2018 and directed EUR 758 million toward technology subsidies. Technology subsidies were, on average, 0.7% of machinery and equipment investment in Finland. This paper is the first quantitative evaluation of the program.

EU Context The program is part of the European Structural and Investment Funds (ESIFs), one of the world's largest industrial policy programs. ESIFs aim to support economic development across all EU countries, especially in remote regions. The 2014–2020 program budget was EUR 670 billion.[12] The national government and the EU fund technology subsidies together, typically 50/50. Decisions are made locally by the ELY Centers. The EU regulates the budget and rules for giving subsidies. The study speaks to the firm-level effects of the broader EU program.

The Program's Objectives The technology subsidies aim to promote the adoption of new technologies. The agenda behind this objective is to improve firms' competitiveness. Technology subsidies in Finland have a long tradition based on the idea that the government can foster growth and structural change through industrial and regional policy (Rodrik 2007; Kekkonen 1952; Mitrunen 2021). The program follows the EU's technology neutrality principle—firms can choose their technology as long as it is new—and is not primarily about the direction of technology, e.g., automation vs. non-automation (Acemoglu 2002a).[13]

A Typical Case A typical technology subsidy is a EUR 100K cash grant paid toward technology costs. The technology is typically a new CNC machine, often combined with a robot, software, or measurement device. The firms are typically SMEs that manufacture fabricated metal products, e.g., parts for large industrial machinery. The subsidies provide funding for up to 35% of the investment, typically 15%. ELY Center pays the grant against verifiable technology costs. Subsidies of this size are audited, and approximately 30% of all ELY subsidies are audited.

The Selection Process The selection process works in three stages, illustrated in Figure 1-4.

1. Application. Starting from all firms, some firms apply for technology subsidies. For our research design, it means that we compare firms that all plan a technology investment. Firms do not apply because a) they do not plan to invest, b) they do not know about the program, c) anticipate they are not eligible, or d) consider the opportunity cost higher than benefits.

[12]Source: ESI Funds Open Data Platform.

[13]The standard economic rationales for the subsidies could be coordination problems, credit and information frictions, and pure transfers to lower-income regions. However, typically in political discourse, the program is not assessed in contrast to the free-market benchmark but seen in the context of economic planning.

2. Pre-screening. In the pre-screening stage, firms contact ELY Centers that pre-screen them before submitting formal applications. This stage is helpful for our design: after pre-screening, the centers' goal is that all firms have a realistic chance of winning the subsidy. The coarse evaluation criteria are size, industry, and general economic position. The program requires the firms to be primarily in manufacturing and SMEs, not owned by large firms, not in financial difficulties and can carry out the technology plan. Firms may decide to skip this stage, but that does not improve their chances of winning the subsidy (but it creates rejected applications from otherwise high-performing firms that are not, e.g., SMEs).

3. Decision. In the decision stage, firms submit a formal application explaining the investment and timeline. Funding is discretionary. Subsidy winners are selected based on the program rules and local and temporal budget priorities and constraints, and an identical firm could receive a subsidy in a given year but not the other. ELY Centers do not score the applications on a formal scale, but we use the evaluation reports to match applicants. In the decision stage, ELY Centers re-evaluate the coarse criteria: size, ownership structure, industry, and financial position. ELY Centers make an impact assessment to evaluate the effectiveness of the subsidy. Cases where the subsidy is more likely to have any impact, are more likely to receive it. Other priorities also exist: firms satisfying the criteria for small firms and firms in remote regions are prioritized.[14] ELY Centers evaluate potential market distortions and sometimes reject applications if the subsidy negatively interferes with local competition. About 15% of applications are rejected.[15]

What Separates Winners from Losers? Text data allows us to read all evaluations of winning and losing applications. Winning applications' evaluations state why the project satisfies the criteria, and the officer recommends a subsidy. Losing applications' evaluations specify why the officer does not recommend a subsidy. Typical reasons for rejection are 1) effectiveness: the subsidy is not expected to affect the project, the project is small and unlikely to have a meaningful effect, the firm had already started the project or received a subsidy for a similar project, 2) industry, size, and investment-type restrictions: the firm is not an SME, e.g., owned by a large firm, a particular industry or investment is not supported at that time or region, the firm proposes to buy used machinery, which is generally not allowed, 3) budget constraints: subsidy funds are limited at that region and time, 4) technical issues: the firm did not provide the required information by the deadline, 5) firm's financial position and the owners' history: ongoing corporate restructuring, foreclosure, or tax liability, and 6) interference with local competition. Employment-related reasons do not appear as typical reasons for rejection; we address this concern in Section 1.7.

Comparing Subsidy Applicants to Average Manufacturers Table A.1 compares the main sample to all Finnish manufacturing firms. Technology adopters are different from non-adopters.

[14] Our regression discontinuity (RD) design is based on changes in the criteria defining a small firm.

[15] Corruption is unlikely to play a significant role in the process. The Corruption Perceptions Index (CPI) ranked Finland as having one of the lowest levels of corruption in 2012–2020.

The subsidy sample firms are larger (despite being SMEs), more productive and profitable, and more educated. Importantly, technology adopters grow faster than average manufacturers. These observations highlight that non-winning applicants provide a better control group than average manufacturers because all applicants have indicated a strong interest in technology adoption. Our estimates capture the local treatment effect for firms close to investing in technologies.

Expected Effects on Technology Investment We conceptualize the technology subsidy as a temporary price reduction for technology. If a firm is close to the margin on whether or not to invest, a temporary price reduction might push it to invest. Firms reported in our interviews that subsidies affect investment because they lower the price of technology, including the associated costs and the future risk of debt. Firms' managers and subsidy officers often mentioned the non-monetary costs of adopting new technology: mental investment and courage. They see the subsidy also as a tool to change the mindset to scale up from an idea to production.

We clarify the source of variation using a model adapted from Cooper et al. (1999) in Appendix A.9. The model maps the price changes induced by the program into the firm's technology adoption decision and factor demand. Under the model, the firm's technology adoption reflects four forces: 1) the replacement cycle, 2) shocks to technologies' prices, 3) shocks to technological progress, and 4) shocks to productivity. Our design based on technology subsidies isolates the role of technology price shocks on technology investment.

1.4.2 Winners-Losers Design

Our main empirical strategy is an event-study design that contrasts similar firms, one of which was approved for technology subsidies while the other was not. The identification strategy is based on the idea that subsidy decisions are quasi-randomly assigned with respect to the counterfactual changes in firm outcomes after conditioning on the information used in the screening process. We assess the comparability of winners and losers and provide several alternative estimation strategies, including a matched non-applicant control group, and matching with text data in the next section.

We estimate two types of equations. Our main specification is the stacked event study:

$$Y_{jt} = \alpha_j + \kappa_t + \sum_{\tau \in \mathcal{T}} \left[I^{\tau}_{jt} \cdot (\gamma_\tau + \beta_\tau \cdot D_j) \right] + X^{\tau}_{jt} + \varepsilon_{jt} \tag{1.1}$$

where Y_{jt} is an outcome for firm j in year t, D_j is the treatment indicator, I^{τ}_{jt} is the event-time indicator for firm j's decision having occurred τ years ago, and the set $\mathcal{T} = \{-5, -4, \ldots, 4, 5\}$ defines the five-year horizon over which we study dynamics. Our parameters of interest are the coefficients β_τ. They summarize the differential trajectory of mean outcomes for winning and losing firms by the time relative to their application. Note that event-time is explicitly defined also for the control group by application year, and firms are only in the treatment or control group for the entire panel.[16] Estimates before the event serve as a test of differential pre-trends between the treatment

[16] Focusing on a control group that never receives treatment reduces the problems arising in the estimation of

and the control group. The coefficients γ_τ capture the common event-time τ effects. The term α_j is the set of firm indicators, κ_t set of calendar-time t indicators, i.e., cohorts of applicant firms, and X_{jt}^{τ} contains potential pre-period controls interacted with both time indicators (the main figures are reported without). We designate $\tau = -3$ as our base event period and omit it. We set the base clearly before the event to avoid contrasting the post-period to any anticipation effects (e.g., Ashenfelter's dip).[17] For clarity, we present all main estimates in reduced form (i.e., intention to treat, ITT).

To summarize the dynamic estimates into a single number, we estimate the stacked first-differences specifications:

$$\Delta Y_j = \beta \cdot D_j + X_j + \varepsilon_j \tag{1.2}$$

where ΔY_j is the change in the outcome from the base year $\tau = -3$ to the post period that we define in each context. The main regressor is D_j, an indicator for whether the firm won the subsidy. We also estimate continuous versions where D_j refers to the amount of subsidies. The control term X_j controls for potential differential trends across firm and application characteristics. We report standard errors that are robust to heteroskedasticity and cluster by firm.

We report the event studies without additional controls. In the first-differences specifications, we control for the baseline firm characteristics at $\tau = -3$ potentially correlated with subsequent changes in our variables of interest: the 2-digit industry and firm size, and calendar-time t fixed effects. We show the results are robust to different controls in the Appendix.

We construct the analysis sample in the following way. We first restrict to technology applications based on the text data. We then restrict to manufacturing and construction industries for three reasons: the program targets these industries, they produce physical outputs, and we have a concrete understanding of what their new technologies are based on our fieldwork.[18] We exclude the largest 5% of applications because they tend to have poor control units. Finally, we restrict to a balanced sample over the five-year horizon.[19] The treatment group is defined by selecting the largest approved subsidy application for each firm. Event-time indicator $\tau = 0$ refers to the year the subsidy application was submitted. The control group is defined by the largest rejected application. Repeated applications for the same project are generally not allowed and untypical.

The ideal experiment that could capture the causal effects of technology on employment, skill demand, and firm performance would randomly assign technology to firms. While a perfect technology experiment is hard to engineer, our identification strategy is based on the quasi-random assignment of technology subsidies, D_j. The identifying assumption is that treatment assignment is conditionally independent of the outcomes:

dynamic treatment effects when the comparison group consists of units that are treated at a different point in time and the event time is not explicitly defined for the control group (Sun and Abraham 2021; Goodman-Bacon 2021).

[17] Our results are robust to the choice of base year.

[18] This leaves out some technology subsidies, for example, for hotels' online reservation systems.

[19] The main reason for this restriction is to ensure that employment and skill estimates come from the same sample; skill shares are only defined for existing firms. We show the results are robust to a non-balanced sample (Table A.14).

Assumption 1 (Rosenbaum and Rubin 1983, CIA): $(Y_{1j}, Y_{0j}) \perp\!\!\!\perp D_j \mid X_j$,

where Y_{1j} and Y_{0j} tell what happens if the firm wins or loses a subsidy.

Our identification strategy exploits the fact that the subsidy program induces quasi-exogenous variation in selection into technology adoption. We compare subsidy-receiving firms to firms that applied for the subsidy but did not receive it. Because the sample includes only pre-screened applicants to the subsidy program, these comparisons control for differences between technology adopters and nonadopters that originate in the decision to apply for technology subsidies. Pre-screened non-winning applicants probably provide a better control group for technology adopters than conventional samples because, like subsidy winners, all applicants have indicated a strong interest in technology adoption. But such comparisons do not control for all criteria used by the program to decide which applicants to accept. The data analyzed here contain information on most characteristics used by the program to accept applicants, including the evaluation report itself (next section). Therefore, the remaining selection bias induced by the decision stage can be eliminated using regression techniques or matching using the information used in the decision process.

Table 1.2 reports summary statistics for the treatment and the control groups. The groups are reasonably similar in terms of revenue, employment, and worker composition. The main differences are that the losing firms are smaller and applied for smaller subsidies. The pre-period differences between the treatment and control motivate our matching strategy in the next section.

An alternative counterfactual is similar firms that did not apply for subsidies. We use coarsened exact matching (CEM; Iacus et al. 2012) to define these similar firms. This matching strategy addresses the concern that the losing firms are not a reasonable counterfactual for what would have happened if the approved firms had not received the subsidy. We match by revenue, employment, wages at $\tau = -3$ plus revenue and employment changes in percentages from $\tau = -3$ to $\tau = -1$ and industries' main sectors (letter classes). The CEM percentiles are 10, 25, 50, 75, 90, and 99. The match is 1:1 with replacement. We define matched control samples for both winning and losing firms; the latter is a placebo test. Tables A.36 and A.37 show the covariate balance for the matched samples. The matched control group also serves to assess whether the patterns in the losing firms are typical or specific to the losing applicants.

1.4.3 Text Matching

We demonstrate a novel method of crafting a research design by controlling for program participants' underlying differences using text data. The subsidy records contain *a report* written by the officer evaluating the application. Given similar reports, treatment assignment is more likely to reflect quasi-random variation than systematic differences. The reports record qualitative characteristics potentially related to the firm's future trajectory. Text matching methods allow us to control for these characteristics (see, e.g., Romer and Romer 2004; Roberts et al. 2020).

As our main text-matching method, we control for propensity scores computed from evaluation reports of applications. The propensity score is a predicted probability that conditional on a text

(W_j), the firm will win a subsidy:

$$p(W_j) \equiv E\left[D_j = 1 | W_j\right]. \tag{1.3}$$

The propensity score theorem (Rosenbaum and Rubin, 1983) states that, in principle, controlling for the probability of treatment allows to satisfy Assumption 1. Propensity scores are valuable in this context as a dimension-reduction tool as directly controlling for texts is not feasible.[20]

The subsidy records contain three types of texts that track the decision process: 1) application summary, 2) evaluation, and 3) decision texts. The application summary and evaluation texts are written by a middle-rank officer responsible for administrating the subsidy and presenting it to a manager for a decision. We use the evaluation texts to compute the propensity scores. These texts capture clearest the potential differences between the firms. Based on our interviews, the subsidy officers' goal is to present an unbiased evaluation.[21]

The text propensity score method works in three steps.

Step 1: We represent the text as data. We use a vector representation based on word embedding. In particular, we employ the FastText (Bojanowski et al. 2016) library for the Finnish language. The advantage of the vector representation is that it captures the semantic meanings of the text instead of a word collection. This is helpful in our context because our goal is to extract information from the evaluations beyond clear markers of success or failure.

Step 2: We estimate the propensity scores using the data. We use a machine learning method, support-vector machines (SVMs), to calibrate the word vectors into probabilities. We train the model on all subsidy applications. The probabilities are calibrated using Platt scaling: a logistic regression on the SVM's scores, fit by five-fold cross-validation on the training data (Zhang, Damerau and Johnson 2002). Figure 1-5 provides the calibration plot for our analysis sample: the predicted probabilities based on text data are on the x-axis and the probability of subsidy receipt on the y-axis. The predicted probabilities closely match the empirical probabilities.[22]

Step 3: We control for confounders using the propensity score. Regression adjustment is our preferred approach. We compare the estimates to coarsened exact matching (CEM) and inverse probability weighting (IPW; Hirano et al. 2003).[23]

As an alternative text-matching method, we use cosine similarity. It measures similarity between two non-zero vectors of an inner product space:

$$\text{cosine similarity } = \frac{\bar{A} \cdot \bar{B}}{\|\bar{A}\| \|\bar{B}\|} = \frac{\sum_{i=1}^{n} A_i B_i}{\sqrt{\sum_{i=1}^{n} A_i^2} \sqrt{\sum_{i=1}^{n} B_i^2}}, \tag{1.4}$$

where A_i and B_i are components of vector $\bar{A}$ and $\bar{B}$. Cosine similarity allows us to compute a

[20] There is only one report for applicant firm j, and hence the propensity score $p(W_j)$ contains only subscript j.

[21] The evaluation text is available for 89% of the main analysis sample.

[22] We calibrate the propensity scores with all possible applications, including exports and R&D. The propensity scores are robust to fully out-of-sample calibration but less precise. We estimate standard errors by bootstrap.

[23] There are multiple ways to implement these steps: represent the text as data, model and estimate $p(W_j)$, and use $p(W_j)$ (Angrist and Pischke, 2009; Gentzkow et al., 2019). The results are broadly robust.

similarity score directly between the texts' vector representations without projecting them first to a single-dimensional propensity score.[24] We construct a matched sample for the winners by selecting the nearest-neighbor with replacement from the losing firms. Table A.2 reports the summary statistics for the cosine-similarity matched sample.

1.5 Estimates

This section provides the reduced-form estimates on employment and wages, skill composition, and firm performance using the primary research design. The main result is clear: we find no evidence of employment reduction or skill bias across a comprehensive set of skills and technologies. The estimates show that after winning a technology subsidy, firms invested sharply more in technologies, hired more workers, but did not change their skill composition. Before receiving a technology subsidy, the winning and losing firms had similar trends in technology investment, employment, and skill composition. The results are robust to controlling for the text propensity score and other controls. The RD and spikes designs in Appendices A.4 and A.3 confirm the results. The results are not limited to the subsidy program or SMEs.

The First Stage Figure 1-6 shows the first-stage event-study estimates β_τ from Equation 1.1. The outcome is technology investment. Winning a subsidy is associated with a sharp increase in technology investment. Before the subsidy application, the groups are on parallel trends. Figure A-1 shows alternative first-stage estimates with all possible subsidies granted and received. It shows that winners and losers are granted a different amount of subsidies exactly in the event year, not before or after. The pattern for received subsidies matches technology investment. Table 1.3 reports the first stage estimates for the main versions of the winners-losers design, with and without text matching. The outcomes are technology subsidies, technology investment, and capital. The first stage is robust to controlling for the text propensity score.

Employment and Wages Figure 1-7 displays the event-study estimates β_τ from Equation 1.1. The outcome is employment relative to the base period $\tau = -3$. The estimates indicate that technology subsidies led to approximately 20% higher employment in the five years after receiving it. As the figure shows, the employment pre-trends were similar between the treatment and control groups. Figure 1-10 visualizes and Table 1.4 reports the first-difference estimates from Equation 1.2, with and without the text propensity control, and with the matched non-applicant control group. These estimates combine the multiple event-study estimates into a single number. Our preferred specification with the propensity control indicates a statistically precise 23% increase in employment. The employment estimates are consistent with the idea that the advanced technologies were a complement to labor in this context.

[24] A conceptual difference is that the propensity score measures the text's predictive power on treatment assignment, while cosine similarity measures the overall similarity between evaluation texts.

Another way of measuring the potential replacement effects of advanced technologies is the labor cost share. It measures the share of revenue that a firm pays to workers. We find a precise zero estimate, reported in Table 1.4. We also generally find a zero effect on wages; in some specifications, there is a small, statistically insignificant negative effect.

The employment estimates are similar when using the matched non-applicant control group (Table 1.4 and Figures A-29, A-31), regression discontinuity design (Figure A-42 and Table A.44), and spikes design without subsidies (Figures A-36, A-38). The employment results are also robust to different text matching versions (Table A.3), different controls (Tables A.4, A.5), and are clearly present in the mean graphs that compare the treatment and control group over time (Figure A-12).

Skill Composition Figure 1-8 displays the event-study estimates for the main firm-level skill measures: average years of education, college-educated workers' share, and the production workers' share. We find no change in these measures, either before or after the technology subsidy. Figure 1-9 summarizes the estimates and Table 1.4 reports the numerical values. Our 95% confidence interval excludes over .15 year changes in the average years of education. The results are in contrast with the view that advanced technologies increase the share of more educated workers and decrease the share of production workers in manufacturing firms. The main skill-composition estimates hold in all our research designs and are robust to a variety of controls referenced in the employment results, including text matching.

We zoom into more detailed skill outcomes: education groups (Figure A-2), occupation groups (Figure A-3), cognitive performance (Figure A-4), school performance (Figure A-5), personality (Figure A-6), demographics (Figure A-7), and task composition (Figure A-8). The big picture is that the effects are primarily skill neutral in the sense that the skill composition does not change. Another central observation is that the baseline skill levels of workers in the sample firms are well below the median. For example, the average cognitive performance is .3 standard deviation lower than the average population, and the average 9th grade GPA is .56 standard deviation below the population average. The sample workers also score lower in tests designed to measure personality traits valued by the Finnish Defence Forces, such as achievement aim and dutifulness. The only personality trait the workers score higher than average is masculinity (+.15 standard deviation). Finally, there are some patterns of changes in the skill composition that are consistent with the observations from our fieldwork, while not statistically significant and subject to multiple testing concerns. The treatment effect on average school GPA is .1 standard deviation (Figure A-5), and the treatment effects on activity-energy, achievement aim, and sociability are .05 standard deviation (Figure A-6). These are the traits the managers and workers we interviewed consistently mentioned to be complementary to new advanced technologies, as opposed to higher education or non-production occupations.[25]

[25]Managers and workers emphasized the non-cognitive skills required: initiative, cooperation, and adaptability, and that workers perform multiple tasks. One CEO explained: "A company does not just pay a welder to weld."

Firm Performance Figure 1-10 visualizes and Table 1.4 reports the first-difference estimates from Equation 1.2 for measures of firm performance: revenue, labor productivity, total factor productivity, and the profit margin. We measure labor productivity as revenue per worker and total factor productivity from Cobb-Douglas production function estimation.[26] The robust finding is that technology subsidies and technology investment led to approximately 30% higher revenue in the five years after. However, we find no evidence of changes in productivity and the profit margin. This potentially surprising finding is consistent with Criscuolo et al. (2019), who study an investment subsidy program in UK manufacturing, and Cerqua and Pellegrini (2014), who focus on capital subsidies in low-performing regions. We provide an interpretation in Section 1.6.

Magnitudes Table 1.5 reports the first-difference estimates from Equation 1.2 with a continuous treatment variable, the subsidy granted in EUR. The estimates from our preferred specification indicate that 1 EUR in subsidies stimulated 1.3 EUR in machinery investment. The firms' revenue increased by 5 EUR per 1 EUR of subsidies.

Table 1.6 reports more detailed estimates on financial outcomes. The average profit margin is 5%. Winning a subsidy led to an increase in average gross profit by EUR 24K and financial costs by EUR 4K. The coefficients from continuous treatment are close to zero. There is a positive .05 effect on financial costs for each subsidy euro granted—that is, the firms carried additional financial costs as a reaction to the subsidy. Because the baseline profitability is moderate in these firms, and they increase their revenue and employment in the same ratio and incur additional costs from the investment, winning a subsidy did not lead to a large increase in profits.

The employment increase is .23 jobs per EUR 10K subsidies, indicating a cost per job of EUR 43K (USD 49K). This number closely matches the numbers managers reported for machinery per worker in their plant in our interviews. Our estimate is close to the average among the cost-per-job estimates reviewed by Criscuolo et al. (2019). It is relatively close to the cost per job estimates of USD 43K by Pellegrini and Muccigrosso (2017) and USD 68K by Cerqua and Pellegrini (2014) in the context of capital subsidies to businesses in the least developed regions in Italy, and the estimate of USD 63K by Glaeser and Gottlieb (2008) for the US Empowerment Zones. Criscuolo et al. (2019) report an estimate of 27K USD at the firm level.

1.6 Mechanism

To recap the results: technology investment led to increases in employment and no changes in skill composition—in contrast with the ideas that technologies replace labor or are skill biased. This section offers a theoretical interpretation and then provides novel theory-motivated tests of that interpretation. We close by explaining when and why we expect to see these results.

[26] TFP is not ideally suited to measure firm performance in our context because (as we will show in Section 1.6) the firms introduce new product varieties. Revenue per worker is robust to different production functions.

1.6.1 Theoretical Framework: Process vs. Product

We outline a framework that contrasts two types of technological change: *process* versus *product*.[27] Process refers to productivity improvements within an output variety, product to the expansion of new varieties. The framework is standard (Dixit and Stiglitz 1977; Melitz 2003; Bustos 2011), but we apply it to a new context. The central element is imperfect substitutability between output varieties. The intuitive distinction is whether firms *use* new technologies to do the same thing at a lower cost or to do new things. We show that these two types of technological change predict different effects and can be empirically tested.

The core idea of the model can be simplified as a composite function:

$$F\left(T_E; f\left(T_I; L\right)\right). \tag{1.5}$$

The function highlights two types of technological change:

T_I **Process** (The Intensive Margin): This affects the production "recipe" f of how factors L are used in production activity. Example: a welding robot replaces human welder's tasks.

T_E **Product** (The Extensive Margin): This affects the "lens" F through which production is projected into markets. Example: a welding robot makes longer seams than a human welder.

1.6.1.1 Setup

Our basic setup is based on Melitz (2003) and Melitz and Redding (2014).[28] The market structure is monopolistic competition with product differentiation and increasing returns to scale at the firm level. The model specifies preference and firm heterogeneity in a differentiated product market. This structure allows technology to have a role in creating new varieties—as in many standard growth models (e.g., Romer 1990). We show that the view of new varieties has different implications than one emphasizing technology's role in allowing productivity improvements within a variety.

Preferences Preferences over sectors $j \in \{0, 1, ..., J\}$ take the Cobb-Douglas form:

$$U = \sum_{j=0}^{J} \beta_j \log Q_j, \quad \sum_{j=0}^{J} \beta_j = 1, \beta_j \geq 0. \tag{1.6}$$

[27] We use the terms process vs. product, but other terms could convey the same idea: e.g., cost vs. differentiation (Porter 1985), secondary vs. primary (Saint-Paul 2002), or defensive vs. enterprise (e.g., Boone 2000). The critical distinction is whether technological change affects how the output is made versus how the customer receives it.

[28] We aim to introduce the simplest model necessary to explain the findings, which captures the essence of a broad class of models featuring process vs. product type technological changes. The Melitz (2003) framework allows for a simple way of introducing imperfect substitutability between varieties. We specifically build on the version by Melitz and Redding (2014). Related approaches include Hopenhayn (1992), Ericson and Pakes (1995), Klette and Kortum (2004), Acemoglu et al. (2018), Akcigit and Kerr (2018), and Hemous and Olsen (2021).

There is a continuum of differentiated varieties within each $j \geq 1$ sector, and these preferences take the Constant Elasticity of Substitution (CES) Dixit and Stiglitz (1977) form:[29]

$$Q_j = \left[\int_{\omega \in \Omega_j} q_j \left(\omega \right)^{(\sigma_j - 1)/\sigma_j} d\omega \right]^{\sigma_j/(\sigma_j - 1)}, \quad \sigma_j > 1, j \geq 1. \tag{1.7}$$

Sector $j = 0$ is a homogeneous numeraire good with a unit-input requirement for production.

The upper-tier Cobb-Douglas preferences imply that consumers spend $X_j = \beta_j Y$ in sector j, where Y denotes aggregate income. The lower-tier CES preferences imply that the demand for each differentiated variety within sector j is:

$$q_j(\omega) = A_j p_j(\omega)^{-\sigma_j}, \quad A_j = X_j P_j^{\sigma_j - 1}, \tag{1.8}$$

where P_j is the price index:

$$P_j = \left[\int_{\omega \in \Omega_j} p \left(\omega \right)^{1 - \sigma_j} d\omega \right]^{1/(1 - \sigma_j)}, \tag{1.9}$$

and A_j is a market demand index, determined by sector spending and the price index. There is a continuum of firms; each firm is of measure zero relative to the market, and takes A_j as given.

Production Firms produce varieties using a composite input L_j with unit cost w_j in sector j. The firms choose to supply a distinct differentiated variety. Production has a fixed cost f_j and a constant marginal cost, inversely proportional to productivity φ. The composite input needed to produce q_j units of a variety is:

$$l_j = f_j + \frac{q_j}{\varphi}. \tag{1.10}$$

Equilibrium We focus on the equilibrium within a sector (and drop the sector j subscript for clarity). The firms choose their prices to maximize profits subject to a residual demand curve with constant elasticity σ. The equilibrium price for each variety is a constant mark-up over marginal cost derived from the first-order condition for profit maximization:

$$p(\varphi) = \frac{\sigma}{\sigma - 1} \frac{w}{\varphi}. \tag{1.11}$$

That gives the equilibrium firm revenue:

$$r(\varphi) = A p(\varphi)^{1-\sigma} = A \left(\frac{\sigma - 1}{\sigma} \right)^{\sigma - 1} w^{1-\sigma} \varphi^{\sigma - 1}, \tag{1.12}$$

[29] This representation has two interpretations: 1) consumers demand differentiated consumption goods with "love-for-variety" preferences (e.g., Grossman and Helpman 1991), or 2) final-good firms demand differentiated intermediate inputs, and a greater variety of inputs increases the "division of labor" (e.g., Romer 1987, 1990). Our context is the technology adoption of intermediate-good producing firms that sell their outputs to final-good producing firms.

and the equilibrium firm profit becomes:

$$\pi(\varphi) = \frac{r(\varphi)}{\sigma} - wf = B\varphi^{\sigma-1} - wf, \quad B = \frac{(\sigma-1)^{\sigma-1}}{\sigma^{\sigma}} w^{1-\sigma} A. \tag{1.13}$$

1.6.1.2 Process

Process advances improve firms' productivity within a variety. This is the intensive margin: It allows firms to produce the same thing more efficiently. The change is on the factor-market side.[30]

We introduce the process advances as in Bustos (2011). The firm has a constant marginal cost $1/\varphi$ within a variety. It can adopt a technology T_I that reduces that cost. Figure 1-11 visualizes the idea. This choice is a tradeoff between a fixed cost f_I and a productivity increase to $\iota\varphi$, where $\iota > 1$. The resulting total cost functions are:

$$l = \begin{cases} f + \frac{q}{\varphi} & \text{if } T_I = 0 \\ f + f_I + \frac{q}{\iota\varphi} & \text{if } T_I = 1. \end{cases} \tag{1.14}$$

Process technology adoption is characterized by sorting according to firm productivity: There is a productivity cutoff φ_I^* above which the firm adopts the technology because the adoption choice involves a tradeoff between a fixed cost and a scaled productivity increase.

The predictions are summarized in Table 1.7.[31] Process-type change predicts increases in revenue, productivity, and profit margin. The intuitive idea is that firms with lower marginal costs produce more and earn higher revenues due to the CES demand structure. Lower marginal costs imply higher measured productivity and profits due to the increasing returns to scale. A distinct prediction from the process-type technological change is zero effect on product composition. There is no similarly precise prediction on exports, which depends on whether the exports are new varieties or not. The prediction on prices is negative if the process change is a cost reduction and positive if it is a quality improvement.

The process view nests several standard models of technology and labor.[32] The predictions on employment, labor share, skill composition, and wages depend on the underlying structure of the process change. In the basic setup, firms use a composite factor L to produce the varieties. If that composite factor is only labor, the model predicts a reduction in the labor share as the firm takes wages as given and revenue per input increases. The models where technological change reduces costs and affects labor typically assume that technological change is "skill biased" in the sense that new technologies complement high-skill workers and increase their share of employment. If the technological change is automation (Acemoglu and Restrepo 2018), it replaces tasks performed by

[30] The process efficiency motive is present in the models of specialization (Smith, 1776), labor-saving technologies (Marx, 1867), growth (Solow, 1956), routine-replacement (Autor, Levy and Murnane, 2003), tasks (Acemoglu and Autor, 2011), automation (Acemoglu and Restrepo, 2018), product and process (Utterback and Abernathy, 1975), and in the 'Schumpeterian models' (Grossman and Helpman, 1991; Aghion and Howitt, 1992).

[31] We derive these predictions in Appendix A.8.

[32] For example, the canonical (Tinbergen 1975; Katz and Murphy 1992), routine-replacement (Autor et al. 2003), and automation models (Acemoglu and Restrepo 2018).

labor with capital and reduces the labor share.

1.6.1.3 Product

Product advances enable the production of new varieties. This is the extensive margin: It allows firms to produce new things and switch between varieties. The change is on the product-market side. Critical to this view of technological change is that outputs with different types are imperfect substitutes. In our framework, there is only one dimension to improve productivity, but multiple dimensions to change product attributes. There is only one firm per variety (the most productive), but firms can differentiate through multiple varieties.[33] [34] [35]

We introduce the product advances by adapting from Melitz (2003). The firm can introduce a new variety by adopting a technology T_E. Figure 1-11 visualizes the idea. The technology requires a fixed entry cost f_E. Potential entrants to the new variety, both existing and new firms, face uncertainty about their productivity in the new variety. After the firm pays the entry cost, it observes its productivity φ for the new variety, drawn from a distribution $g(\varphi)$, with cumulative distribution $G(\varphi)$. The firm then decides whether to produce or exit the project. Melitz (2003) shows this decision is characterized by a cutoff productivity φ_E^* where the firm makes zero profits:

$$\pi\left(\varphi_E^*\right) = \frac{r\left(\varphi_E^*\right)}{\sigma} - wf = B\left(\varphi_E^*\right)^{\sigma-1} - wf = 0. \tag{1.15}$$

In equilibrium, the expected ex-ante profits equal zero due to free entry:

$$\int_0^\infty \pi(\varphi)dG(\varphi) = \int_{\varphi_E^*}^\infty \left[B\varphi^{\sigma-1} - wf\right]dG(\varphi) = wf_E. \tag{1.16}$$

We visualize the relationship between profits and productivity in Figure A-50. Firms with $\varphi < \varphi_E^*$ would lose if they produced. They exit the project, receive $\pi(\varphi) = 0$ in that new variety, and cannot recover their entry cost. The subset of the firms that produce and have $\pi(\varphi) > wf_E$ make positive profits after the entry cost.

The predictions from the product-type technological change are different from the process type. As shown in Table 1.7, product-type change predicts an increase in revenue but no changes in productivity and profit margin. The intuitive idea is that the new variety allows the firm to sell more, but its productivity and profit margin are still, on average, the same as before due to the

[33] A new variety has several interpretations: a new product, a quality change not perfectly substitutable with quantity, re-purposing production to respond to changing demand, expansion to new markets, capturing a larger share of the value chain, etc. A new variety may be the same product but with an improved process that provides more reliable scheduling or a faster response time to orders, changing the aspects customers receive.

[34] The expansion of variety in consumer and intermediate goods plays a central role in many theoretical models of growth (Romer 1990; Grossman and Helpman 1991). The product view is closely related to Porter (1985): gaining competitive advantage through a quality-differentiation strategy instead of a cost-leadership strategy.

[35] We distinguish two directions of change: vertical (within the same variety) vs. horizontal (a new, imperfectly substitutable variety). In this class of models, vertical cost reductions or quality improvements within the same variety are essentially equivalent. The reason is that the model assumes perfect substitution between quality and quantity within the same variety. The productivity term φ can be interpreted in terms of costs or within-variety quality; the interpretations are isomorphic to a change in units of account (Kugler and Verhoogen 2012).

free-entry condition. Some new varieties are more profitable, some less.

The next distinct prediction from the product-type technological change is the effect on the product composition. While a new variety does not equal a new product (e.g., it could also be a faster response time), a new product is a signal of a new variety. Exports are also a signal of new varieties. If different markets have differentiated preferences, a new variety makes the firm more likely to export, export a larger share of its revenue, or export to a larger variety of destinations. If the new variety is a quality improvement, the predicted price effect is positive.

The predictions on employment, labor share, labor composition, and wages again depend on the underlying structure of the product change. But this time, the critical difference is that there is no unambiguous basis for expecting a sustained effect on the share or composition of labor. The skill or task composition might differ for a new variety, but that depends on the particular context. However, the basic structure predicts an increase in the use of the composite factor, generally employment (see also Harrison et al. 2014). The model predicts zero wage effects in a competitive labor market (for both technological advances) since wages are determined in the sectoral equilibrium and the firm is small relative to the market.

1.6.2 Evidence: Testing Process vs. Product

This section empirically tests whether the technological changes we observe are the process vs. product type. We document that they are primarily the product type. This observation helps explain the puzzling results of no labor replacement or skill bias. Firms used new technologies to create new types of output, not to replace workers.

We proceed in two steps. First, we directly measure the type of technological changes using our text and survey data. Second, guided by the framework, we consider a new set of outcomes that are critical signals that contrast process vs. product type change.

1.6.2.1 Directly Measuring the Type of Technological Change

We measure the type of technological change directly using text and survey data.

Text Data Text data allow us to read the sample firms' technology adoption plans. Based on our theoretical framework, we code the technology projects into process vs. product. Process refers to using technologies to produce the same type of output more efficiently, while product refers to using technologies to produce a new type of output or expand.

Figure 1-12 shows that 91% of projects in our sample are of the product type. These applications describe new products, access to new markets, responding to changing demand conditions, growth, or similar use for the technology. Only 9% of the texts do not describe such reasons. The technological changes we document are primarily product advances based on this measure, and our sample contains few purely process-type technological advances.[36]

[36]Our interviews suggest that while process-type advances exist, they are less likely to be physical machinery but new management styles such as lean manufacturing and digitization.

While the sample is mostly product type, we estimate treatment effects separately for the two categories. We use the matched control group described in Section 1.4.2 because our control sample is small for both categories. Table A.8 provides some evidence that product advances led to larger employment effects and no skill bias. Process advances led to smaller employment effects and some skill bias, .14 years, significant at the 10% level.

Survey Data We also measure the uses of technologies with survey data. The European Community Innovation Survey (CIS) asks our sample firms and other firms about the importance of different objectives for process and product innovations. The options include introducing a more extensive product selection, quality improvement, and lower labor costs.

Figure 1-13a shows that typical reasons for firms' process and product innovations are access to new markets, introducing a larger product selection, better quality, and larger capacity. Lower labor costs rank the 6th most important: only 20% of firms report that lowering labor costs is important for process and product innovation. Based on CIS data, we code the firm's technology project as the product type if the firm considers one of the product-type reasons (in black) important but does not consider lower labor costs important. Conversely, we code the technology project as the process type if lower labor costs (in grey) are important, but none of the product reasons are. Figure 1-13b shows that 97% of our technology-adoption cases are the product type. These numbers are similar when considering our spikes design sample, all manufacturing firms, or all Finnish firms, suggesting that the finding is not limited to the subsidy program. Our interviews with CEOs corroborate the observation from the survey data.[37]

1.6.2.2 Testing the Predictions with New Outcomes

Process and product type technological change predict different effects, summarized in Table 1.7. We use these predictions to distinguish them. So far, we have shown that the technological advances—either with or without the subsidies—led to increases in employment and revenue, no change in skill composition, the labor share, wages, productivity, or the profit margin. These empirical results are consistent with the product-type predictions but not with the process type. Next, we provide evidence for new outcomes: exports, products, marketing, prices, and patents, all signals of product-type changes.

Figure 1-14 shows the event-study estimates with exporter indicator as the outcome. Subsidy winners are more likely to become exporters. Table 1.8 reports a treatment effect of 4 percentage points from the baseline of 28%. The effect on the exports' revenue share is .9 p.p. from the baseline of 5.2%. The winners also start exporting to .2 more regions, from 1.5 baseline.[38]

Table 1.8 reports the effects on products, measured from the customs data. The treatment effect is .15 products from the baseline of 1.55. We also observe an increase in the product *turnover*: the

[37] Table A.9 shows the estimates by the technology category using the survey data. We use a matched control group since the original control group's overlap with the survey is limited. The estimates for the product group are similar to the overall group. The process group is too small to estimate the results (marked by –).

[38] The export results are consistent with Lileeva and Trefler (2010) and Koch et al. (2021).

treatment firms both introduce and discontinue more products.

Figure 1-15 shows that subsidy winners are more likely to increase their marketing expenditure. The increased marketing signals that the firms intend to change how the customers perceive their output—a product-type change—not only their production costs.

Table 1.9 reports the treatment effects on prices. We measure prices from the Customs Register and the Industrial Production Statistics (a survey of manufacturing firms). We focus on product-level prices' unweighted average. We find a 29.1% increase in the customs data prices and 30.8% in the manufacturing survey. Price increases signal potential quality improvements.

Figure A-11 shows the evolution of the subsidy applicant firms' patenting status. While suggestive evidence, we observe that patenting is concentrated in the periods before applying for subsidies and technology investment. This pattern of patenting is an additional signal that firms used the subsidies and technologies to scale up from an idea to production.

Some research proposes that exports and new products are also skill biased (Bernard and Jensen 1997; Xiang 2005; Matsuyama 2007). One reason we do not observe skill bias from exports or new products is that these changes—which we conceptualize as product advances—are a normal part of how these firms operate. We observe in our fieldwork that these manufacturers constantly identify shifts in demand and redeploy their productive resources to new uses using new technologies. Also the large-scale manufacturers combine economies of scale with flexibility, reflected in short production runs, product introductions, and sensitivity to customer needs. Earlier fieldwork by Dertouzos et al. (1989), Berger (2013), and Berger (2020) corroborates these observations.

1.6.3 Two Types of Manufacturing: Mass Production vs. Flexible Specialization

Our theoretical framework tells a tale of two types of technological change—process vs. product—and how they predict different effects that can be empirically distinguished. A central question created by our empirical analysis is: when and why is one more likely to occur than another? The technology adoption events in our data are almost entirely product rather than process-type changes. But both types may occur in reality, and some studies report examples of the latter when it comes to automation (e.g., Acemoglu and Restrepo 2020; Restrepo and Hubmer 2021). We explain next why our findings are distinctive but logical—and applicable to other settings where similar incentives for process vs. product type technology adoption prevail.

To do so, we contrast two types of manufacturing: mass production (Taylor 1911; Ford 1922) vs. flexible specialization (Piore and Sabel 1984; Milgrom and Roberts 1990). These two different *contexts* affect the incentives for the two types of technological change. Mass production is characterized by standardized products, high volumes, and a stable environment, and it makes process advances more likely. Flexible specialization is characterized by specialized products, low volumes, and an unstable environment. It makes product advances more likely.

Our results differ from the two views emphasized in the literature—that technologies replace labor or are skill biased—because the literature has focused more on process advances in mass production (e.g., Acemoglu and Restrepo 2018). In contrast, the flexible manufacturing system is more common among the firms we study. In our context, both small and large manufacturing firms produce specialized products in small batches. Examples include defense contractors building specialized equipment and industrial manufacturing firms producing new wind power stations. However, the findings may not apply to the mass production of non-specialized commodities, such as cement or steel, or high-volume assembly, where costs are critical.

A large literature documents that manufacturing has moved from mass production to new, more flexible, and specialized forms of production since the 1980s (e.g., Dertouzos et al. 1989; Berger 2013). These new forms of production emphasize quality and responsiveness to market conditions while utilizing technologically advanced equipment. Piore and Sabel (1984) call this change the second industrial divide, Kenney and Florida (1993) call it moving beyond mass production, and Milgrom and Roberts (1990) call it modern manufacturing. While different studies approach the topic from different angles, the common observation is that "the business environment is no longer conducive to producing standardized products for a stable market" (Piore 1994). One of the managers in Berger (2020) explained clearly: "American manufacturing has been transformed. It's become highly engineered, highly specialized, and highly customized. I see this across all manufacturing. This is a different country. It's no longer the mass production of the past." Why did this change happen? The research suggests several reasons: consumers shifted away from standardized goods (Bils and Klenow 2001), globalization reduced the cost of specialization between firms (Berger and Center 2005), and new technologies reduced setup times and made it less costly to switch production between products (Bartel et al. 2007).

Next, we help understand when and why process vs. product type technological advances are more likely, and how this trade-off relates to the type of manufacturing—mass production vs. flexible specialization. We point out three central factors: scope for specialization, volume, and the need for adaptation that each affect the incentives for process vs. product type changes.[39]

Specialization The trade-off between process versus product advances depends on the scope for specialization. Firms in a sector with a higher scope for specialization are more likely to implement product advances, and a lower scope for specialization makes process advances more likely (see also Sutton 1998; Kugler and Verhoogen 2012). Intuitively, in sectors with a higher scope for specialization, firms may gain a competitive advantage by introducing a new good or changing their selection of goods. This contrasts with sectors that produce bulk goods, where the primary source of competitive advantage is cost. Scope for specialization comes most naturally in the framework from the elasticity of substitution σ_j in sector j: A higher elasticity magnifies the effects of productivity improvements on revenue and profitability (Appendix A.8). The intuition is that when the elasticity of substitution is high, demand is more responsive to price reductions,

[39]These are not the only factors that may influence the choice. Other relevant factors include: automation feasibility (Graetz and Michaels 2018; Acemoglu and Restrepo 2020), employment protection (Saint-Paul 2002; Manera and Uccioli 2021), complementary resources, such as venture capital, trade associations, and suppliers (Berger 2013; Gruber and Johnson 2019), and skill supply (Dertouzos et al. 1989; Berger 2013).

making process advances that reduce costs relatively more effective.

One measure to capture the scope for specialization is the Rauch (1999) index based on whether the good is a commodity.[40] Figure 1-16 shows that 91% of the firms are in an industry with a Rauch index over .5, indicating a high scope for specialization. Our main industries, fabricated metal products, machinery and equipment, and wood products, have an index of 1 and are fully specialized based on the Rauch index. Our sample does not include firms in non-specialized industries, such as cement, steel, or paper.[41] Specialized manufacturing is not limited to the subsidies design: the share of firms (and employees) in specialized vs. non-specialized industries is similar in the spikes design and Finnish manufacturing overall.

Table A.11 reports further evidence: the number of firms by the scope for specialization and technology category. Less than 1% of our sample are process advances in non-specialized industries (e.g., cost reductions in steel manufacturing or automation in the paper industry). Consistent with our interpretation, product-type projects are more common in specialized sectors.[42]

Volume The trade-off between process vs. product depends on the production volume. In our interviews, most managers explained that they are specialized low-volume producers who invest in advanced technologies to make the products they sell to a few customers with unique demands. Our theoretical framework rationalizes why technology adoption events are more likely to be the product than process type in a low-volume context. In the framework, the amount of input required to produce volume q_j of a variety is:

$$l_j = f_j + \frac{q_j}{\varphi}, \tag{1.17}$$

where f is the fixed production cost and $1/\varphi$ is the constant marginal cost. The process-type technology adoption decision T_I is a tradeoff between an additional fixed cost f_I and a productivity increase to $\iota\varphi$. The high-volume producers benefit more from the productivity increase because the fixed cost is distributed over the higher volume. The low-volume producers benefit less from the productivity increase, but not from the introduction of new products. In our model, high-volume firms are also large firms with low marginal costs because, given the CES demand structure, firms' relative outputs and revenues inversely depend on their relative marginal costs.

Looking at the evidence, firms in our sample are mainly SMEs, as shown in Table 1.2, consistent with observing mainly product-type technology adoption events. Tables A.6 and A.7 describe the matched product and process samples. The groups are similar because our context is relatively uniform, but there are some relevant differences. Consistent with our interpretation, the product-type firms are smaller.

[40] Measures of the scope for specialization also include Gollop and Monahan (1991) and Sutton (1998).

[41] Dertouzos et al. (1989) emphasize that even in steel manufacturing, quality improvements are crucial.

[42] Table A.10 provides treatment-effect estimates for specialized vs. non-specialized industries. The estimates are generally similar in both groups. Our interpretation is that because the clear pattern in our data is product-type technological change in specialized industries, it is unsurprising that we do not observe different effects in the small subsample of firms in the non-specialized industries.

Adaptation Over time, the trade-off between process vs. product depends on the need for adaptation. Most firms we interviewed described operating in a changing environment where adaptability is important. One manufacturer described they could automate their assembly—currently done manually—but it would require them to commit to a specific model and set of parts to build it. This commitment was unattractive as they must update their model and parts frequently to stay competitive for their customers. In this context, the firm had more substantial incentives to use technologies to create new varieties than to improve its productivity within a variety. This need for adaptation arises from, for example, changes in consumer preferences, technological obsolescence, and cost competition. A firm we interviewed explained: "We cannot compete with the low-cost competitors. We need to offer unique goods and services."[43]

We conceptualize the need for adaptation as a death shock that occurs with an increasing probability $\delta \in (0,1)$, adapted from Melitz (2003):

$$\delta \in (0,1), \quad \frac{\partial \delta}{\partial t} > 0. \tag{1.18}$$

The death shock increases the relative incentives for the product-type technology. It generates a discount factor for the value computation and reduces the net present value of future revenue in the given variety and, therefore, reduces the benefits from the process-type technological change. In contrast, with a new variety, the firm can start with a lower death risk.

Our text data directly records that firms invest in technologies to respond to changing demand. The need for adaptation also has two key empirical predictions: 1) we will observe a higher product turnover in addition to new products, and 2) we observe a negative trajectory for those firms that did not adopt the technology and a higher survival for those firms that did. Our evidence confirms both predictions (Table 1.8, Figures A-12, A-13).

1.7 Robustness

We conduct several robustness checks to evaluate the internal and external validity of our findings.

1.7.1 Internal Validity

Selection Bias A natural concern when estimating the impact of technology adoption is the bias due to a potential correlation between the adoption and unobserved characteristics of adopters. These concerns are less likely to be important in our setting because (as described in Section 1.4) we focus on variation induced by a technology subsidy program, where comparisons by adopter status are restricted to a sample of applicants to the program. Non-adopting applicants probably provide a better control group for adopters than conventional cross-section samples because, like

[43] Firms with limited capabilities to respond to cost competition may launch new varieties when faced with low-cost rivals (Porter, 1985; Aghion et al., 2005). This idea is consistent with Bloom et al. (2016) and Fieler and Harrison (2018), who document that import competition induced innovation and product differentiation. Bernard et al. (2010) analyze product switching as a source of reallocation within firms.

adopters, applicants have indicated a strong interest in technology adoption. Moreover, the data analyzed here contain information on most characteristics used by the subsidy program to screen applications. The selection bias induced by subsidy program screening can therefore be eliminated using regression techniques or by matching on the covariates used in the screening process. Our results are robust to controlling for the pre-application characteristics and the evaluation report texts (Tables 1.4, 1.5, A.3, A.4, and A.5).

To directly investigate whether the rejected applications are a reasonable counterfactual for the approved applications, we read through all approved and rejected applications in the analysis sample. We found only ten rejected applications that did not seem likely to receive subsidies in any situation: either the entrepreneur had a concerning history or the firm's financial position was unstable. Our results are robust to excluding these applications. We also find similar effects when using a matched non-applicant control group (Appendix A.2). As a placebo test, we contrast the main control group to a matched non-applicant control group. We find no first stage on investment and a small positive transitory effect on employment, indicating that the subsidy losers grew somewhat faster than similar non-applicant firms.

We use three different research designs: 1) the winner-losers design, 2) a regression discontinuity design using unanticipated changes in the subsidy program rules (Appendix A.4), and 3) an event-study design focusing on technology adoption events (Appendix A.3). These designs generate similar results. This suggests that selection bias in any single design is unlikely to drive our results.

The remaining concern is selection bias common to all our research designs. The concern would be that none of the control groups we analyze here represents a reasonable counterfactual for technology adopters. To address this concern, we can analyze trends in adopter firms without any control group. Figure A-12 shows the evolution of treatment group means for machinery investment, employment, and years of education. Machinery investment increased sharply after the technology subsidy application; winners increased their employment but did not change their skill composition disproportionately. Trends in technology adopters do not support the view that advanced technologies reduced employment or significantly changed skill composition.

Statistical Power A concern particularly relevant to presenting a null result is statistical power. Are our results precise and technology-adoption events large enough to justify our conclusion about no significant changes in skill composition measured by education and occupation? The estimates from our preferred specification indicate a -.004 change in the average years of education at the firm level, with a standard error of .075 years, meaning that we can exclude over .15 year increases in the average education. In comparison, the treatment and control firms increase their education on average over the 5-year event window by .4 years.

The small effects could be driven by small events. Several aspects suggest that this is not the reason for our findings: 1) A typical technology adoption event in the subsidy sample is EUR 100K, a doubled investment compared to an average year. The monetary value is a lower bound: the purchase price of the machinery is only part of the total cost, about 25% in the US manufacturing documented by Berger (2020). The rest of the cost is the machine bed, installation, and all the work needed to integrate the machinery into the plant. 2) The subsidy program requires that the technology investments represent significant technological advances to the firm. 3) We consider large technology investment events in the spikes design in Appendix A.3 and find null effects on skill composition measured by education and occupation.

1.7.2 External Validity

There are several legitimate external validity concerns and alternative explanations for our findings and interpretation. To repeat here: we do not argue that our results apply everywhere. We document typical technological advances in manufacturing firms in Northern Europe. While we acknowledge that other technological advances exist, our fieldwork suggests we do not document a marginal phenomenon. Next, we respond to specific external validity concerns.

Concern 1: The subsidy program is biased toward employment and low-skill work. The observation behind this concern is, to some degree, correct. One of the objectives of the ELY Center subsidy program is to stimulate employment by supporting the adoption of advanced technologies in manufacturing firms. But several aspects support the view that the program's biases are not the primary source of our findings: 1) We find similar results also when evaluating technology adoption events without the subsidy program. 2) Interviews with managers document that the subsidy-supported technology adoption events are not notably different from typical technology adoption events. 3) Interviews with subsidy administrators document that significant technology projects are unlikely to be rejected because they would not stimulate positive employment effects.[44] 4) To address this concern systematically, we read all rejected applications and investigated whether they were rejected for employment-related reasons. In none of the applications was the concern about employment the main reason. Five reports mentioned employment, but the concerns were primarily about the potentially low first stage on technology investment; employment was secondary. Our findings are robust to excluding these applications. Text records also uncover that ELY Centers often interpret the employment effects compared to the counterfactual where the firm is not competitive in the market without the technology and would need to reduce employment; maintaining employment is seen as an increase. 5) The employment effects are not enforced: the firms are free to make their employment decisions after receiving the subsidy. 6) We have no evidence that the program intends to increase low-skill jobs; in fact, ELY Centers support hiring high-skill workers into manufacturing firms.

Concern 2: Workers are already skilled and learn new skills. This alternative explanation proposes that since workers are already skilled and learn new skills, we do not observe changes in skill composition even if technologies are skill biased. To some degree, this is true. Most workers

[44]Some insignificant technology projects get rejected because they are insignificant and unlikely to stimulate technological advances in production and employment effects.

in our sample have specialized training in production work and regularly participate in continuing vocational training (CVTS Survey 2015). All managers we interviewed reported that they combine technology adoption with worker training. New manufacturing technologies require new skills, but our observations from the field indicate that production workers are best suited to learn to use them. At the same time, the debate on skill bias has focused on the idea that advanced technologies replace production work and increase the relative demand for college-educated workers. We do not find evidence of either at the firm level.

Concern 3: The technologies are not typical advanced manufacturing technologies. A natural concern is that our estimates capture something other than the effects of standard advanced technologies in manufacturing, particularly that we miss the effects of automated technologies. To address this concern, we classify technologies into automated versus non-automated technologies using text and customs data, as described in Section 1.3. Automated technologies are considered automated in everyday language: e.g., robots, CNC machines, and conveyor belts. Non-automated are manually operated: e.g., non-automatic welding tools, hydraulic presses, and cutting machines. In our text data, non-automated refers to all applications not classified as automated. Figures A-9 and A-10 show the estimates of firm-level effects for automated vs. non-automated technologies. The effects are similar in both groups, and we still find employment increases and no changes in the skill composition from automated technologies. Finally, the spikes design captures major technology investment events in the industry and size range. While there may be different types of technology adoption events, our estimates capture the average of these events.

Concern 4: Credit constraints drive the employment and skill effects. One alternative explanation is that the effects are primarily about access to credit rather than technologies (an exclusion restriction concern). While credit constraints are likely to play a role in allowing the subsidies to induce firms to invest more, several arguments work against this explanation for the employment increases and skill null result: 1) We observe a strong first stage on technology investment. 2) We do not observe larger effects for the ex-ante more likely credit-constrained firms: small firms (Table A.12) and firms with higher financial costs (Table A.13). 3) We observe the same effects without the program in Appendix A.3.

Concern 5: Fixed costs in production lead to skill neutrality. One concern is that these firms could have non-homothetic production technologies where fixed and variable costs have different factor intensities (Flam and Helpman 1987). The fixed costs could be educated managers and technical staff, while the variable costs could be production workers. If the firms use technologies to expand, the increase in variable costs could mask the potential skill bias of technologies. This concern has a testable implication: it should be less important for large firms. Small firms might primarily increase their variable costs, while we would expect that large firms would also need to scale their fixed costs. Table A.12 reports the main estimates by firm size. We find no significant differences, suggesting that non-homothetic production is unlikely to be the cause for our findings.

Concern 6: Firm-level employment gains replace employment elsewhere. A firm's technology adoption may affect other firms, and the total employment and skill effects may differ from those reported here. Two aspects make estimating these effects challenging: 1) the firms are relatively small, and 2) they trade globally directly or indirectly through their customers; thus, externalities are likely to be minor. Theoretically, whether or not the technology adoption events replace employment elsewhere depends on the type of technology and the kind of externalities it induces. We document that our technological advances are the product type: the firms use technologies to produce new output types. These outputs are typically intermediate goods or machinery for final-good producing firms. In Romer (1990), this type of variety expansion generates growth—that is, some of the externalities may be positive. At the same time, new intermediate goods could replace previous vintages of intermediate goods as in the "Schumpeterian models" with quality improvements and creative destruction as in Grossman and Helpman (1991) and Aghion and Howitt (1992). Exploring these channels is a promising avenue for future research.[45]

1.8 Conclusion

This paper provides novel evidence on a classic question: What are the effects of advanced technologies on employment and skill demand? Our paper is the first to evaluate advanced manufacturing technologies' effects using a research design based on direct policy variation. Our novel administrative data allow us to measure firms' technology investment and workers' employment, wages, and skills precisely over time. To address external validity, we evaluate technology adoption events also without the program.

Our main finding is that advanced technologies, such as CNC machines, welding robots, and laser cutters, did not reduce employment, replace production workers, or increase the share of highly educated workers in industrial and custom manufacturing firms. We find that these technologies led to increases in employment and no change in skill composition. The findings are consistent across all estimation methods, with and without the subsidy program.

This paper proposes a simple explanation for the findings. We document that the firms used new technologies to produce new types of output, not replace workers with technologies. Direct evidence shows that technology adoption led to more revenue, new products, and new exports. Text analysis of firms' technology-adoption plans shows that they adopted new technologies to introduce new products, access new markets, respond to changing demand, and grow. To explain our findings, we outline a theoretical framework that contrasts two types of technological change: process versus product (e.g., Utterback and Abernathy 1975; Porter 1985). Process change refers to productivity improvements within an output variety; product expanding to new varieties (e.g., Dixit and Stiglitz 1977; Melitz 2003). Our evidence indicates that firms invested in advanced technologies to gain a competitive advantage by introducing new varieties. For example, the piston manufacturer we observed invested in new technologies to manufacture more effective pistons.

[45] Acemoglu et al. (2020b), Koch et al. (2021), and Oberfield and Raval (2021) analyze potential externalities.

The results stand in contrast with the view that new technologies reduce employment or increase the share of highly educated workers in manufacturing firms. While no single study can be decisive, we review a body of evidence indicating that technology investments in manufacturing led to increases in employment and to no detectable changes in skill composition (e.g., Doms et al. 1997; Koch et al. 2021).

We do not argue that our results apply everywhere. We obtain our findings in a context where small and large manufacturing firms produce specialized products in small lot sizes. But the findings may not apply to non-specialized commodities, such as cement or steel, or high-volume assembly, where prices and costs are critical. Our results differ from the two views emphasized in the literature because it has focused more on process advances in mass production (e.g., Acemoglu and Restrepo 2018). In contrast, the flexible manufacturing system is more common among the firms we study. Qualitative evidence documents that a large part of manufacturing has evolved from mass production (Taylor, 1911; Ford, 1922) to flexible specialization (Piore and Sabel, 1984; Milgrom and Roberts, 1990). Currently, a large part of manufacturing is specialized.

Our results do not directly apply to non-physical technological advances, such as ICT or the internet (e.g., Autor et al. 2003; Akerman et al. 2015; Gaggl and Wright 2017), management practices, R&D, technological advances in offices, historical eras, or the future. Some technological advances have also replaced workers (e.g., Acemoglu and Restrepo 2020; Bessen et al. 2020), and our results do not challenge the view that skills and technologies are related (e.g., Lewis 2011). Our evidence from the field suggests that work and skill requirements change in subtle ways due to technology investment (as in Bartel et al. 2007).

Our results provide new evidence on the effects of one type of industrial policy: a lump-sum transfer to increase technology adoption in manufacturing firms (see also Criscuolo et al. 2019). Several researchers argue that lack of access to financial support limits the manufacturing sector's ability to scale up ideas into production (Dertouzos et al., 1989; Berger, 2013; Gruber and Johnson, 2019). We find that it is possible to stimulate technology investments by targeted subsidies and, by doing so, induce increases in employment, revenue, exports, and product variety.

Finally, our study makes some methodological contributions. We demonstrate novel methods to use text data in program evaluation. Many policy programs leave a trail of text records, and these texts allow measuring things that would otherwise be difficult to measure. We show how to use text data to measure variables of interest and perform matching. In the spirit of Roberts et al. (2020) and Mozer et al. (2020), we demonstrate how to craft a research design by controlling for program participants' underlying differences using text data. As new technologies have proliferated across firms, so, too, has the empirical literature on their effects. In light of the results reported here, some more conventional estimates of the effect of technologies in manufacturing firms do not appear to be too far off the mark (e.g., Doms et al. 1997).

Main Figures and Tables

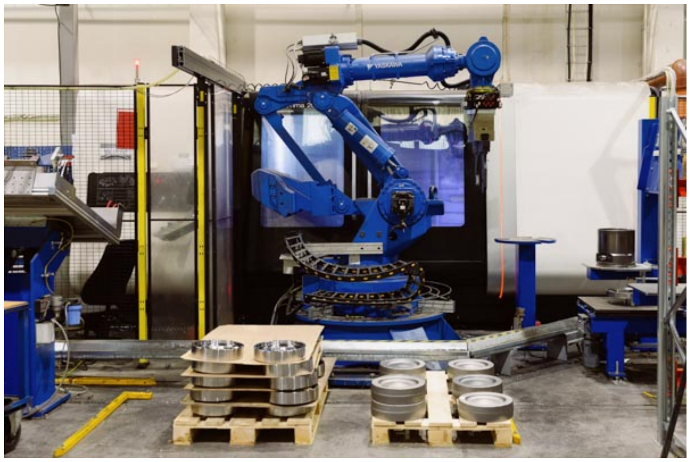

(a) CNC Machine and a Robot.

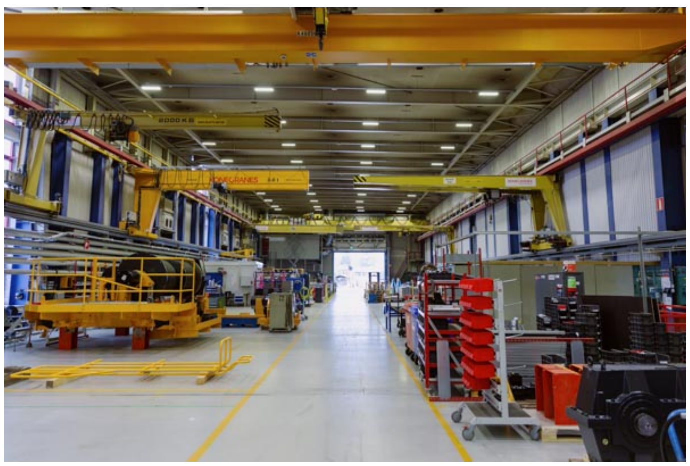

(b) Inside an Industrial Manufacturing Plant.

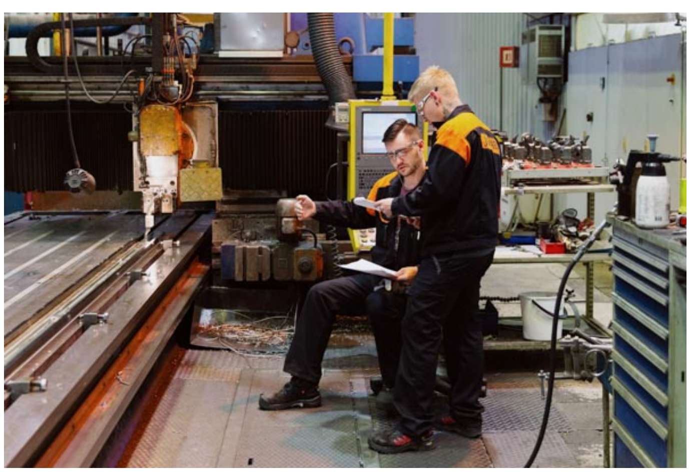

(c) Machine Operators and a Milling Machine.

Figure 1-1: Fieldwork: Documenting the Context.

Back to Section 1.2.

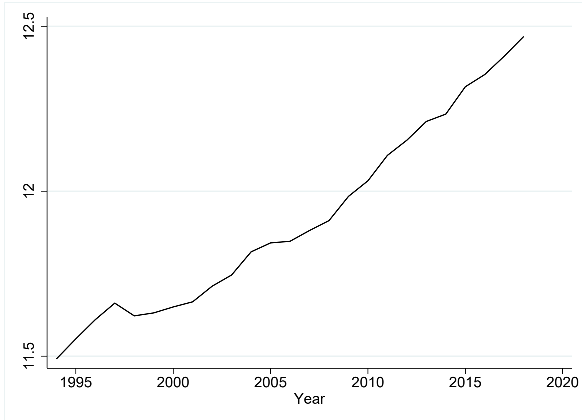

(a) Average Years of Education.

(b) Production Worker Employment Share.

(c) College vs. Non-College Wage Ratio.

Figure 1-2: Manufacturing Skill Trends.

Notes: These figures document trends in Finnish manufacturing over 1994–2018. We restrict to firms with at least 3 workers. We compute the year-level averages from firm-level observations. The numbers are unweighted to match our research design. The employment-weighted numbers are similar. Back to Section 1.2.

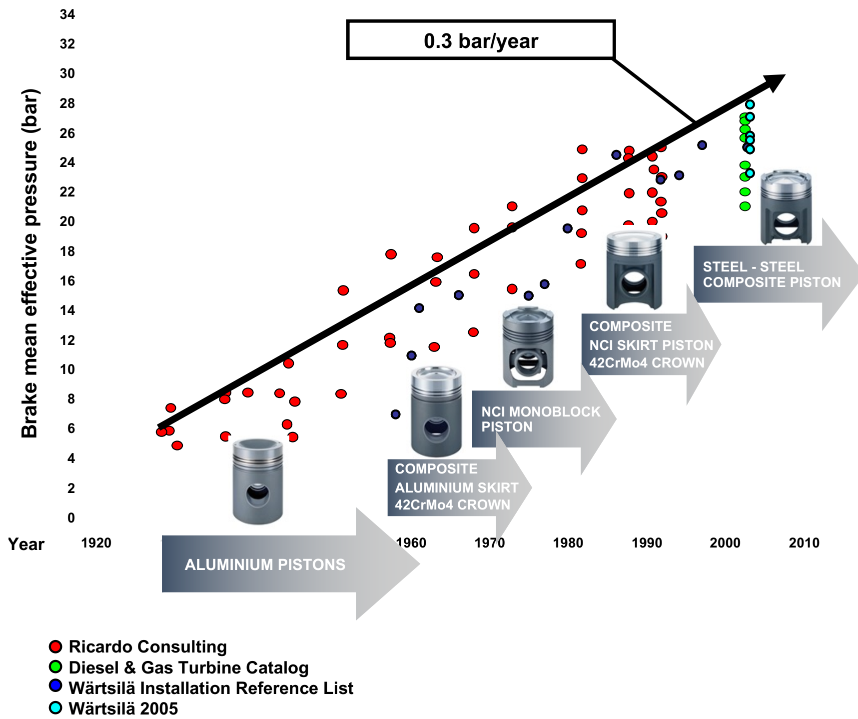

Figure 1-3: Moore's Law for Pistons: The Development Trend of Piston Materials Over 100 Years.

Back to Section 1.2.

Figure 1-4: The Subsidy Application Process.

Notes: Details in the main text. Back to Section 1.4.

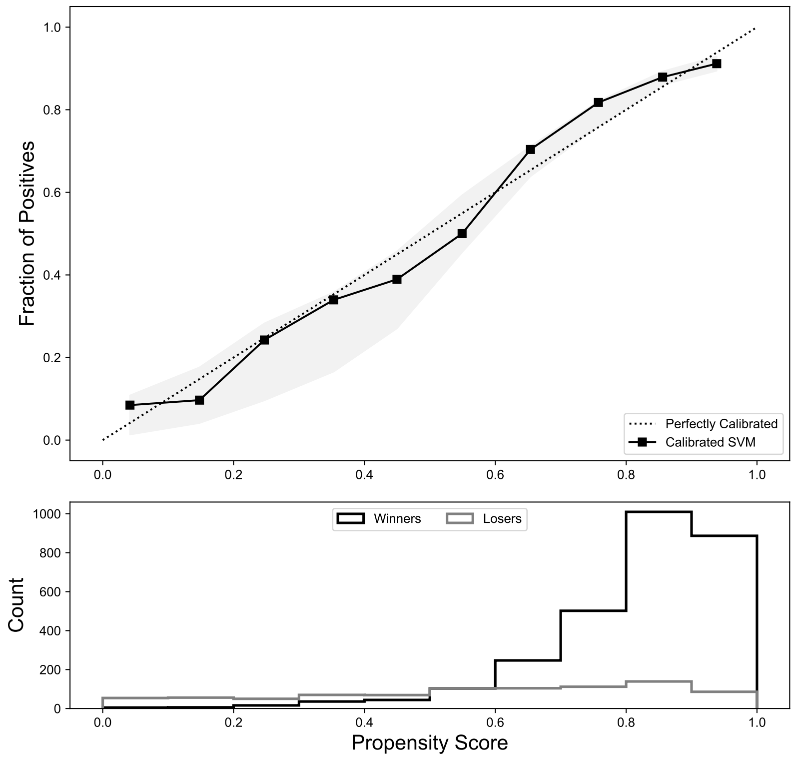

Figure 1-5: The Text Propensity Score Calibration Plot.

Notes: Upper panel: The predicted probabilities of subsidy receipt based on text data are on the x-axis, and the observed probabilities are on the y-axis. The text data are evaluation reports of the applications written by the subsidy program officers. The predicted probabilities are calibrated using a vector representation of the text and SVM. Standard errors are estimated by bootstrap. The predicted probabilities closely match the empirical probabilities. Lower panel: Distribution of the predicted values. Most of the applications have high predicted values reflecting the overall acceptance rate. Back to Section 1.4.3.

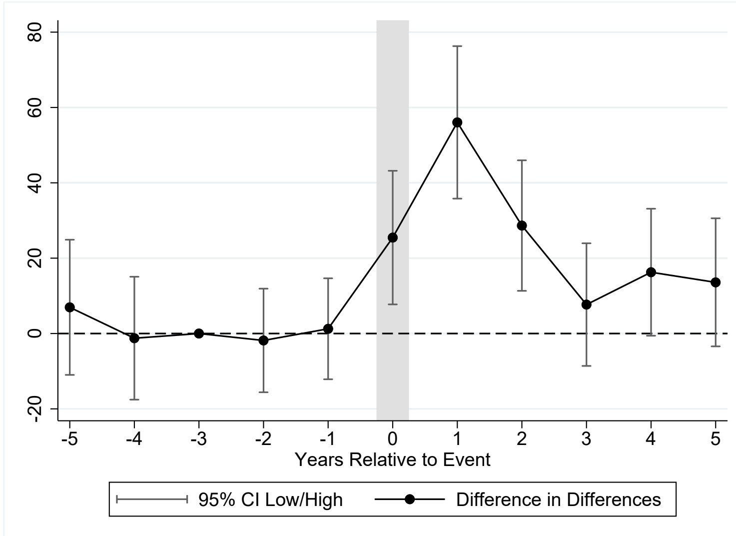

Figure 1-6: The First Stage: The Effect of Technology Subsidies on Machinery Investments.

Notes: Event-study estimates from Equation 1.1. The outcome is investment in machinery and equipment (in EUR 1000s) measured from the financial statement register. Event time $\tau = 0$ refers to the application year. The estimate for $\tau = 1$ indicates that the treatment group invested EUR 60K more than the control group in the year after subsidy application. The estimates indicate a cumulative EUR 130K effect on machinery investment. This event-study specification contains no controls in the term X_{jt}^{τ} of Equation 1.1. Back to Section 1.5.

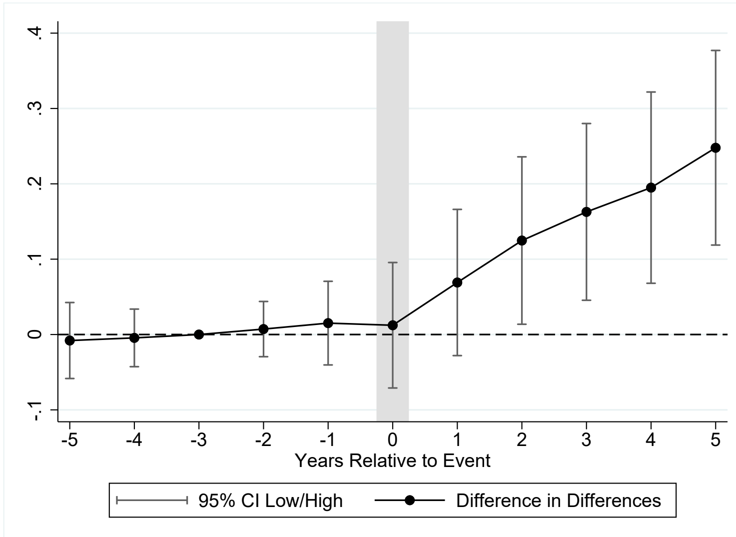

Figure 1-7: Employment Effects: The Effect of Technology Subsidies on Employment (in %).

Notes: Event-study estimates from Equation 1.1. The outcome is employment relative to the base year $\tau = -3$. Event time $\tau = 0$ refers to the application year. The estimates indicate approx. 20% increase in employment. This event-study specification contains no controls in the term X_{jt}^{τ} of Equation 1.1. Back to Section 1.5.

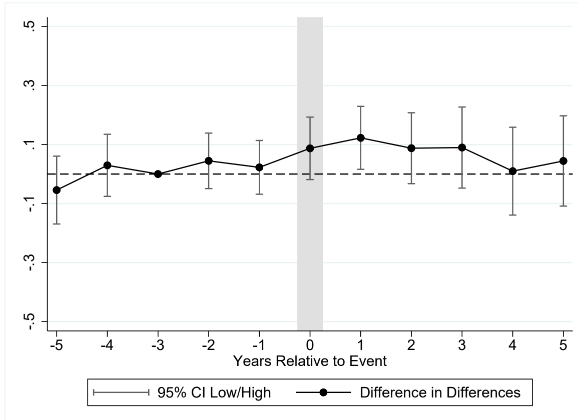

(a) Education Years.

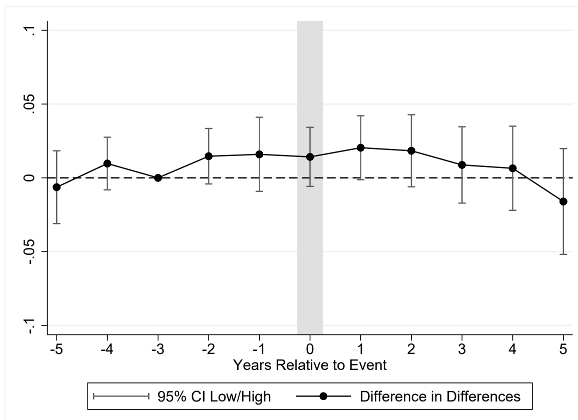

(b) College-Educated Workers' Share.

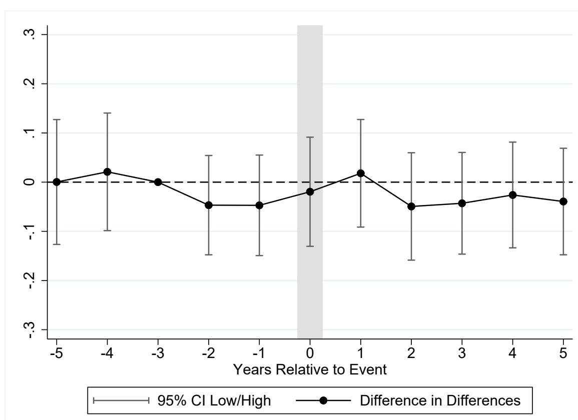

(c) Production Workers' Share.

Figure 1-8: Skill Effects: Event-Study Estimates.

Notes: Event-study estimates from Equation 1.1. The outcomes are relative to the base year $\tau = -3$. Event time $\tau = 0$ refers to the application year. The estimates indicate approximately zero changes in the main skill measures. Education years are defined as the average years of education among the workers in the firm (measured in years); college-educated workers' and production workers' shares are the shares of employment of that group (measured in percentage points). These event-study specifications contain no controls in the term X_{jt}^{τ} of Equation 1.1.
Back to Section 1.5.

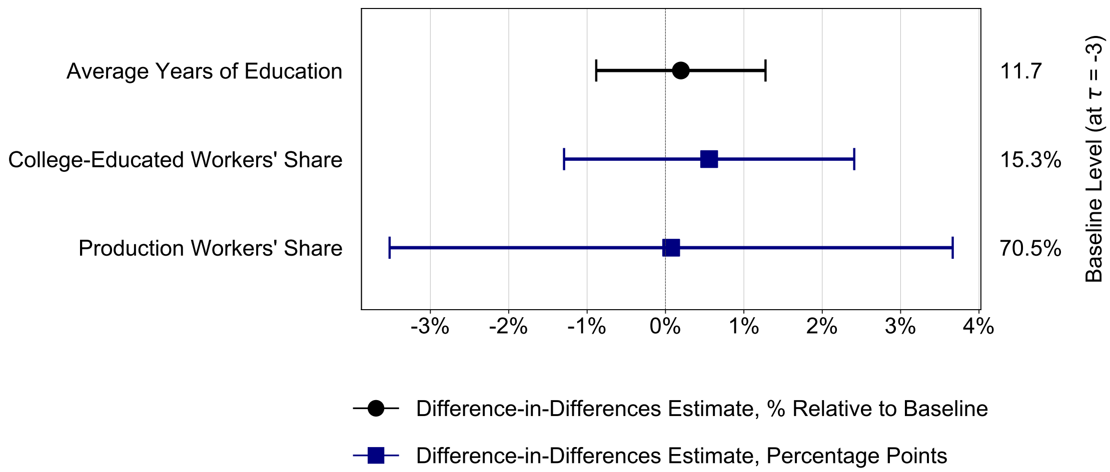

Figure 1-9: Skill Effects. The First-Difference Estimates.

Notes: Difference-in-differences estimates from Equation 1.2. The right-hand side reports means at $\tau = -3$. Education is measured as a relative change (%) in the average years of education in the firm between $\tau = -3$ and the average of $\tau \in [2, 5]$. The shares are measured in percentage-point changes. The estimates indicate no detectable changes in the skill composition. The specifications include two-digit industry and firm size as controls. Back to Section 1.5.

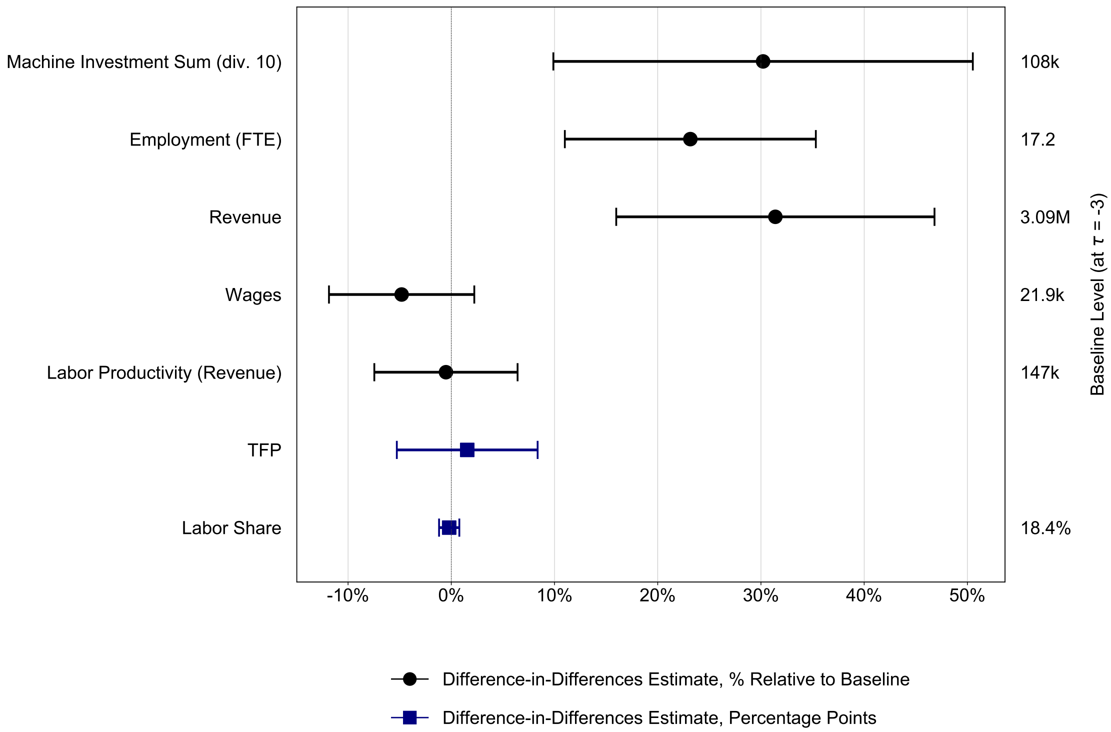

Figure 1-10: Firm-Level Effects.

Notes: Difference-in-differences estimates from Equation 1.2. The right-hand side reports means at $\tau = -3$. Machine Investment, Employment, Revenue, Wages, and Productivity are measured by relative changes to baseline at $\tau = -3$. For Machine Investment, the post-period outcome is the sum of investment between $\tau \in [0, 2]$ and for other outcomes, the average of $\tau \in [2, 5]$. The specifications include two-digit industry and firm size as controls. Back to Section 1.5.

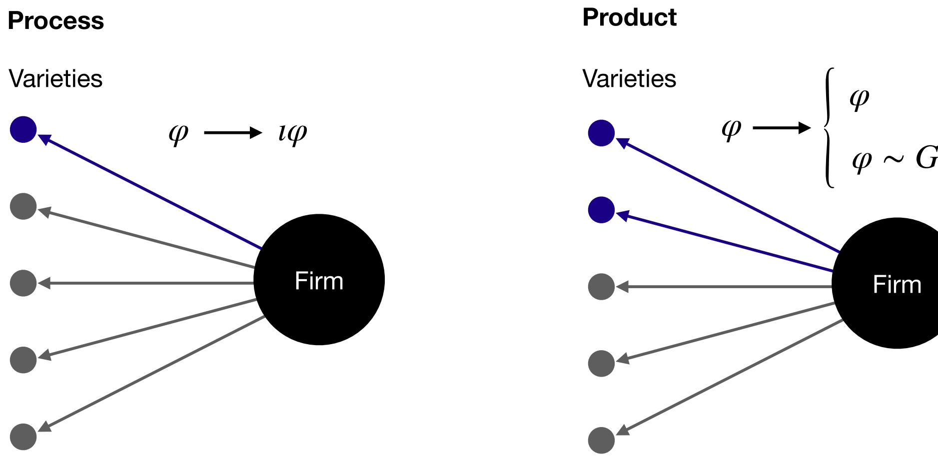

Figure 1-11: Process vs. Product.

Notes: Process refers to productivity improvements within an output variety, product to the expansion of new varieties. Details in the main text. Back to Section 1.6.

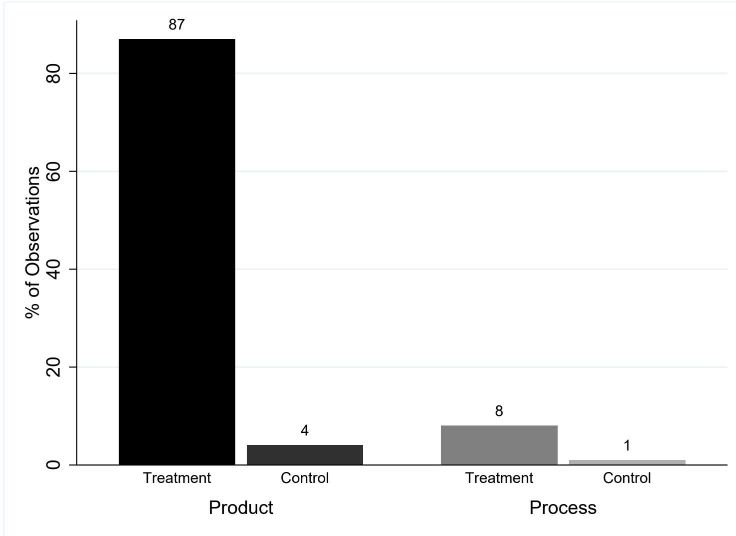

Figure 1-12: Technology Categories Measured from Text Data: Observations by Category.

Notes: Product refers to technology projects that aim to produce a new type of output. Process refers to technology projects that aim to produce the same type of output. The text data are text records from the subsidy program's administration, including each firm's application and evaluation texts. A trained panel performed the classification. Details in the main text. Back to Section 1.6.2.

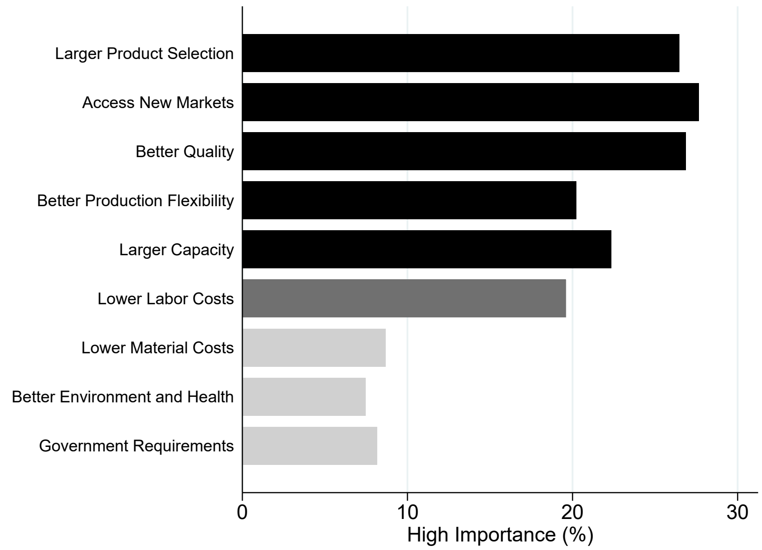

(a) Specific Objectives.

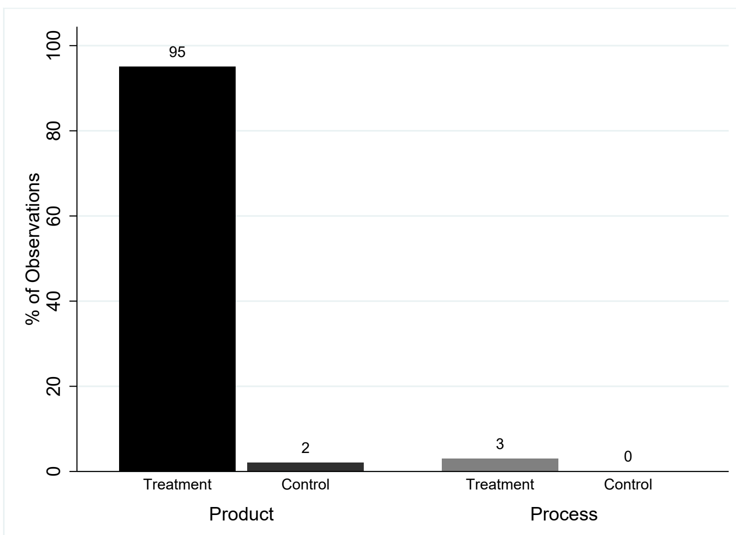

(b) Aggregated Objectives.

Figure 1-13: Technology Categories Measured from the Survey Data: Observations by Category.

Notes: The European Community Innovation Survey (CIS) reports firms' views on the importance of different objectives for process and product innovations, including technology adoption. **Panel (a)** shows the share of firms in our main sample that report the objective is highly important. Variables are in thematic order (new varieties, expansion, costs, environment, and regulations). We use survey years 1996–2008. If the firm has responded to multiple rounds of CIS, we consider the closest survey to its technology-adoption event. **Panel (b):** Product refers to firms that reported that one of the first five objectives was important and lower labor costs were not. Process refers to firms that reported that lower labor costs were important but did not report any of the first five objectives as important. N = 510 (i.e., the number of main-sample firms also in CIS). Back to Section 1.6.2.

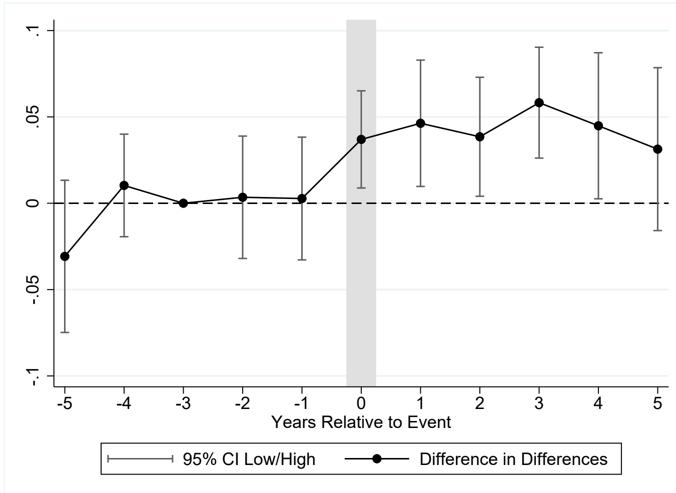

Figure 1-14: Export Effects: The Export Status.

Notes: Event-study estimates from Equation 1.1. Event time $\tau = 0$ refers to the application year. The outcome is the firm's export status indicator (exporter vs. non-exporter). Exports are measured from the Finnish Customs' Foreign Trade Statistics. Export status is measured using the definition by Statistics Finland. A firm is defined as an exporter in a given year if its total export value is over EUR 12K during the calendar year spread over at least two different months, or a single export event is over EUR 120K in value. This event-study specification contains no controls in the term X_{jt}^{τ} of Equation 1.1. Back to Section 1.6.2.

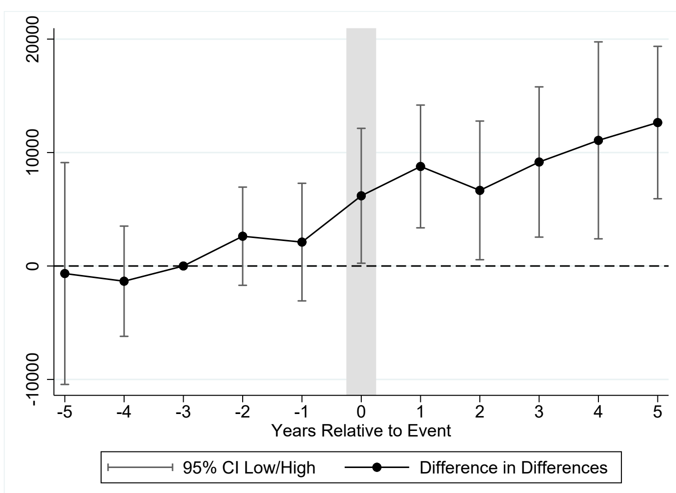

Figure 1-15: Marketing Effects: Marketing Expenditure.

Notes: Event-study estimates from Equation 1.1. The outcome is the firm's marketing expenditure, measured from the Finnish Financial Statement Register. Event time $\tau = 0$ refers to the application year. This event-study specification contains no controls in the term X_{jt}^{τ} of Equation 1.1. Back to Section 1.6.2.

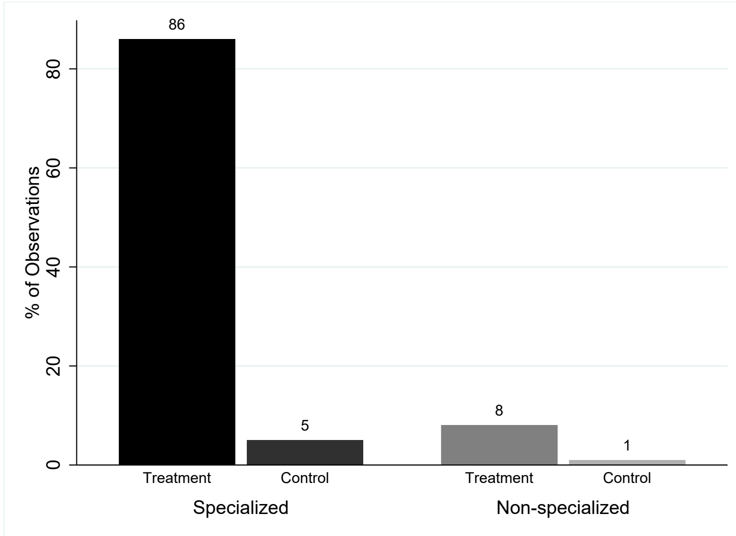

Figure 1-16: Specialized vs. Non-Specialized Industries: Observations by Category.

Notes: Specialized refers to industries producing non-commodities and non-specialized refers to industries producing commodities measured by the Rauch (1999) index. The distribution is similar when using Gollop and Monahan (1991) and Sutton (1998) indices. Back to Section 1.6.3.

Table 1.1: Technology Categories.

Classification	Description
Technologies	All technology investments and projects.
Uses of Technologies	
Process	Produce the same type of output using technologies.
Product	Produce a new type of output using technologies.
Types of Technologies	
Automated vs. non-automated	Technologies with no active user vs. an active user.
Hardware and/or software	Physical vs. non-physical technologies.

Notes: Technologies are measured from the financial, text, customs, and survey data. Uses of technologies are measured from the text data of the technology subsidy program and from the Community Innovation Survey (CIS). Types of technologies are measured from the text data and the customs data. The technology classes are described in Appendix A.5. Back to Sections 1.3 and 1.6.2.

Table 1.2: Summary Statistics: The Main Research Design (Winners vs. Losers).

	Treatment Group		**Control Group**		**Both**		
Variable	Mean	Std. Dev.	Mean	Std. Dev.	10p	Median	90p
Machinery Inv. (EUR K)	109.93	369.14	82.60	233.11	0.00	27.24	233.80
Revenue (EUR M)	3.20	25.39	1.64	5.29	0.16	0.96	5.67
Employment	17.81	47.16	9.67	21.29	1.40	7.90	37.00
Wages (EUR K)	22.23	9.08	18.40	10.22	11.26	22.30	31.61
Subsidy Applied (EUR K)	112.05	129.25	47.01	81.30	8.89	58.13	290.06
Subsidy Granted (EUR K)	81.77	103.02	0.00	0.00	3.24	35.64	200.23
Educ. Years	11.71	0.99	11.45	1.12	10.50	11.73	12.67
College Share (%)	15.51	16.80	11.63	18.42	0.00	12.50	33.33
Production Worker Share (%)	70.53	21.53	70.37	28.61	42.86	72.73	100.00
Observations	1885		146		2031		

Notes: All variables measured at $\tau = -3$. Machinery investment is measured from the financial statement register. Data on revenue, employment, and wages come from the firm- and worker-level registers. Subsidies applied and granted are from the subsidy application data. Education years, college share, and production worker share are measured based on the worker composition within the firm. Back to Section 1.4.2.

Table 1.3: The First Stage.

	(1) Granted Subsidy		(2) Machine Inv. (EUR K)		(3) Capital Stock (EUR K)	
Treatment	66.06***	70.22***	107.9***	100.4***	49.78**	41.60
	(3.119)	(4.907)	(17.53)	(21.90)	(18.26)	(23.60)
Propensity Score		✓		✓		✓
Observations	2031	1812	2031	1812	1560	1540

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2 with and without the text propensity control. To measure capital, we use the official records on firms' balance sheets. The post-period outcomes are sums between $\tau \in [0, 2]$. The specifications include two-digit industry and firm size as controls. Back to Section 1.5.

Table 1.4: Firm-Level Effects.

Panel A: Investment, Employment, and Revenue.

	Machine Investment (EUR K)			Employment			Revenue		
	Baseline	Prop. Score	Match	Baseline	Prop. Score	Match	Baseline	Prop. Score	Match
Treatment	107.9^{***}	100.3^{***}	127.9^{***}	0.232^{***}	0.234^{**}	0.217^{***}	0.314^{***}	0.333^{***}	0.261^{***}
	(17.53)	(21.90)	(6.556)	(0.0614)	(0.0746)	(0.0183)	(0.0779)	(0.0958)	(0.0232)
Observations	2031	1812	3200	2031	1812	3200	2031	1812	3200

Panel B: Wages, Profit Margin, and Productivity.

	Wages			Profit Margin			Productivity		
	Baseline	Prop. Score	Match	Baseline	Prop. Score	Match	Baseline	Prop. Score	Match
Treatment	-0.0481	-0.0285	0.00306	0.00121	-0.00791	-0.00685^{*}	-0.00516	-0.00622	0.0117
	(0.0355)	(0.0407)	(0.00290)	(0.00772)	(0.00978)	(0.00290)	(0.0350)	(0.0427)	(0.0120)
Observations	1952	1738	3080	2031	1812	3200	2031	1812	3200

Panel C: Labor Share and Skill Composition.

	Labor Share			Education Years			College Share			Production Worker Share		
	Baseline	Prop. Score	Match	Baseline	Prop. Score	Match	Baseline	Prop. Score	Match	Baseline	Prop. Score	Match
Treatment	-0.00202	-0.000700	-0.00293	0.0246	-0.00385	0.0303	0.00557	0.00592	0.00542	0.000723	-0.0213	-0.00464
	(0.00496)	(0.00601)	(0.00203)	(0.0611)	(0.0752)	(0.0207)	(0.00935)	(0.0116)	(0.00330)	(0.0181)	(0.0212)	(0.00605)
Observations	2031	1812	3200	1884	1676	2999	1884	1676	2999	1891	1692	3011

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2. The table reports the treatment effects on selected outcomes for the main sample with and without the text propensity-score control and the matched control sample. "Baseline" refers to a baseline specification with calendar-year indicators, two-digit industry, and firm size as controls. "Prop. Score" refers to estimation with the text propensity score included as a control. "Match" refers to estimation in the matched sample, where the control group is formed from matched non-applicant firms. **Panel A:** Machine investment is in EUR K. Employment and revenue are in relative changes, e.g., 0.20 would refer to a 20% increase. **Panel B:** Wages and productivity are relative changes; the profit margin is in percentage points. **Panel C:** Education years is in years. The labor, college, and production worker shares are in percentage points. For machinery investment, the post-period outcome is the sum of investment between $\tau \in [0, 2]$ and for other outcomes, the average of $\tau \in [2, 5]$. Back to Section 1.5.

Table 1.5: Continuous Treatment Estimates.

	(1) Machine Inv.		(2) Employment		(3) Revenue	
Granted Subsidy	1.321***	1.262***	0.249***	0.230***	5.292***	4.973***
	(0.0806)	(0.0809)	(0.0213)	(0.0220)	(0.468)	(0.478)
Propensity Score		✓		✓		✓
Observations	2031	1812	2031	1812	2031	1812

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2. Treatment is the subsidy amount in EUR, scaled to EUR 10K for employment. For machinery investment, the post-period outcome is the sum of investment between $\tau \in [0, 2]$ and for other outcomes, the average of $\tau \in [2, 5]$. The specifications include two-digit industry and firm size as controls. Back to Section 1.5.

Table 1.6: The Effects on Profits and Financial Costs.

Panel A: Win/Lose.

	(1) Profit Margin (%)	(2) Gross Profits	(3) Net Profits	(4) Fin. Costs
Treatment	0.121	24.49*	20.35*	4.133**
	(0.772)	(9.941)	(10.09)	(1.425)
Baseline	5.2	274.0	-16.07	290.1
N	2031	2031	2031	2031

Panel B: Continuous Treatment.

	(1) Gross Profits	(2) Net Profits	(3) Financial Costs
Granted Subsidy	-0.0353	-0.0878	0.0525***
	(0.0638)	(0.0646)	(0.00949)
Baseline	274,006	-16,074	290,080
N	2031	2031	2031

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The effects on profits and financial costs. The baseline means are measured at $\tau = -3$. The profit margin is measured in percentage points. Gross and net profits refer to profits before and after financial costs. **Panel A**: The treatment is the win-lose status. The profits and financial costs are measured in EUR 1000s. **Panel B**: The treatment is the amount of subsidies the firm was granted. The coefficients are interpreted as the effect of one euro in subsidies on profits or financial costs, measured in euros. The baseline medians are 5.0% (profit margin), EUR 52K (gross profits), EUR 37K (net profits), and EUR 8.3K (financial costs). The specifications include two-digit industry and firm size as controls. Back to Section 1.5.

Table 1.7: Predictions from Process vs. Product Type Technological Changes.

Outcome	Process	Product
Revenue	↑	↑
Productivity	↑	0
Profit margin	↑	0
Products	0	↑
Export status and share	–	↑
Employment	–	↑
Labor share	↓	–
Skill composition	↑	–
Prices	↓ if cost	0
	↑ if quality	↑ if quality

Notes: Details in the main text. The symbol – refers to no clear prediction. Back to Section 1.6.1.

Table 1.8: Export and Product Effects.

	(1) Export Status	(2) Export Share	(3) Export Regions	(4) Products	(5) Products Introduced	(6) Products Discontinued
Treatment	0.0404**	0.00935*	0.219***	0.155**	0.0880**	0.0664**
	(0.0134)	(0.00451)	(0.0568)	(0.0599)	(0.0282)	(0.0223)
Baseline	0.284	0.0523	1.498	1.546	0.498	0.539
N	2031	2031	2031	2031	2031	2031

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2 for the main research design (winners vs. losers). Exports and products are measured from the Finnish Customs' Foreign Trade Statistics. Export status is measured using the definition by Statistics Finland. A firm is defined as an exporter in a given year if its total export value is over EUR 12K during the calendar year spread over at least two different months, or a single export event is over 120K EUR in value. The specifications include two-digit industry and firm size as controls. Back to Section 1.6.2.

Table 1.9: Price Effects.

	(1) Price (Exports)	(2) Price (Manufacturing)
Treatment	0.291	0.308**
	(0.328)	(0.102)
N	400	217

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2 for the main research design (winners vs. losers). We winsorize price data at the 10% level within product and year. Prices are measured as product-level revenue divided by quantity from the Finnish Customs' Foreign Trade Statistics and the Industrial Production Statistics (a survey of manufacturing firms). The specifications include two-digit industry and firm size as controls. Back to Section 1.6.2.

Chapter 2

Psychological Traits and Adaptation in the Labor Market

WITH RAMIN IZADI

2.1 Introduction

Economic research documents that negative labor-market shocks, such as unexpected job loss or the disappearance of manufacturing work, cause long-lasting adverse effects on workers (Jacobson et al., 1993; Autor et al., 2014). However, some adapt better than others.[1] In particular, a recent literature demonstrates the predictive power of psychological traits in the labor market (Deming, 2017; Jokela et al., 2017; Edin et al., 2021). But little is still known about the role played by psychological traits in adaptation in the labor market.

This paper provides novel evidence on the significance of psychological traits in adapting to mass layoffs and plant closures. How do personality and cognitive characteristics help workers recover from economic changes? To answer this question, we construct measures of cognitive ability, extraversion, and conscientiousness by applying exploratory factor analysis to classified data from the Finnish Defence Forces. These data contain results from a standardized personality and cognitive ability tests administered to 79% of Finnish men born between 1962 and 1979 (n = 489,252).[2] We combine the military data with the register data of Statistics Finland on employment, wages, education, occupation, and firm performance. Our main empirical results analyze mass layoffs and plant closures in 1995–2010 and estimate the heterogeneous treatment effects with

[1]For example, a line of research shows that the magnitudes of the negative effects can depend on family background and socioeconomic status (Hoynes et al., 2012; Kaila et al., 2021).

[2]The three distinct factor variables are allowed to correlate with each other (pairwise correlation between all three is about 0.4). Conscientiousness and extraversion belong to the so-called Big Five personality taxonomy. Each of the five traits is associated with a group of subtraits or facets. Our underlying test data were not designed with the Big Five model in mind but do include many of the facets as test items. Conveniently, our factor analysis groups the test items approximately along the theoretical lines. Namely, outward-oriented items, such as sociability, leadership ability, activity-energy, and confidence, load onto one factor (which we label "extraversion"), whereas inward-oriented items, such as deliberation and dutifulness, load onto another factor (which we label "conscientiousness").

respect to the measures of cognitive ability and personality.

To set the stage, we document the baseline impact of a mass layoff event on individuals' labor-market outcomes. We include all workers of the firm prior to the event in the analysis to allow the selection into job loss to be part of adaptive behavior. Consistent with literature, we find long-lasting negative effects on earnings. But in contrast to research that finds mostly transitory negative effects on employment, in our sample, employment effects also seem to persist, even if decrease over time (Schmieder et al., 2018; Lachowska et al., 2020).

Our novel main estimates interact the treatment with psychological factors in an event-study framework. This allows us to see how the treatment effects change for different psychological profiles. We estimate the interactions jointly in a saturated regression to account for the cross-correlation of the factors. For each factor, we find a distinct pattern in relation to the treatment effect. Conditional on other traits, extraversion is the only trait that predicts better recovery even in the long term. A one standard deviation increase in extraversion predicts a 20% smaller earnings loss each year. The effect lasts for at least eight years after the event. For the first years after the shock, high cognitive ability also reduces the earnings loss by 20%, but this boost is short-lived and fades out after a few years. In contrast, conscientious individuals do no better or worse than the average individual. We repeat these estimations using employment as the outcome and find that, across traits, the patterns in the reductions of dis-employment are similar to those of earnings.

To understand the drivers of personality's adaptive value, we analyze the potentially adaptive behaviors, such as changing occupation and industry and re-education. Workers who experience a mass layoff event are also much more likely to change occupation or industry. However, we find that psychological traits have relatively little predictive power on these margins of adaptation. If anything, extraverted individuals change occupations and industries less than the average individual. Extraversion predicts faster re-employment in the same type of job rather than re-allocation to a different type of job.

One key question arising from our results concerns selection. To what extent do our findings just reflect differential pre-layoff selection into occupations, industries, and education? Each of these choices can independently influence adaptation and are likely to be endogenous to earlier-life psychological traits. For example, due to occupational and educational selection, extraverts could face less tight labor markets after the shock. To address this, we estimate our main specification with education, occupation, and industry controls. We find that the addition of controls reduces the estimate for cognitive ability significantly but does not influence the estimate for extraversion much. Moreover, extraversion seems to be a better predictor of recovery than years of education. In summary, occupational or educational selection are not the likely drivers of the positive effects of extraversion.

Since we study all individuals who were employed in the downsizing establishments, we can study differential retention rates across traits. Are extraverted or high cognitive ability individuals more likely to retain their employment in a mass layoff? We find that in the long term, high cognitive ability individuals are no more likely to remain in the establishments relative to the average individual, but they are one percentage point more likely to be "early leavers." In other words, they leave the establishments just before the event. At the same time, compared to the baseline exit rate of 50%, this effect is small. Extraverts, on the other hand, are two percentage points less likely to leave the establishment relative to the average. This retention effect persists in the long term and appears to be partly driven by selection into different occupations and tasks within the firm.

Overall, extraversion and cognitive ability predict smaller scarring effects of mass layoffs by helping particularly the extraverted to either keep their jobs or find work more quickly once they are laid off. Of course, in the spirit of heterogeneity analysis, this predictive effect should not be interpreted as a causal effect of extraversion.

This paper brings together two active lines of economics literature: (1) the importance of psychological traits in the labor market and (2) the impact of job loss on workers' outcomes. Importantly, it also re-visits an earlier primarily theoretical literature on adaptation.

Adaptation. Classic theoretical research in economics (Nelson and Phelps, 1966; Welch, 1970; Schultz, 1975) emphasizes the value of skills not just applied to production tasks but adapting to "disequilibria" or changing economic conditions. Empirically, little is known about these adaptation processes. It is unknown how specific skills and traits, such as personality traits and cognitive abilities, influence the adjustments to major economic changes. Our paper combines this classic question in economics with novel psychological measurement. For example, Schultz (1975) leaves it as an open question of whether the skills needed for adaptation are rooted in education or psychological traits. Our analysis shows that particularly extraversion helps workers adapt more than education does, even when controlling for selection into occupations and industry.

Psychological Traits. A large literature analyzes the role of noncognitive skills in the labor market. The evidence unambiguously demonstrates that a wide array of noncognitive skills—personality traits, interpersonal skills, and other features—are important drivers of labor-market success (e.g., Heckman et al. 2006; Lindqvist and Vestman 2011; Deming 2017; Jokela et al. 2017).[3] One limitation is that these existing results consider labor-market outcomes overall in the cross-section. Our paper contributes to understanding the importance of these psychological traits specifically under times of change. An open question is whether the same skills that help people achieve higher earnings also help adapt and recover from shocks. We show that while returns to education, cognitive ability, and conscientiousness are large in the cross-section, extraversion predicts adaptation better. This suggests that one mechanism that makes extraversion important could be related to its value in times of change. The adaptive value of extraversion could be an important source of its overall value in the labor market. Conversely, conscientiousness does not predict resilience to labor-market shocks in our context.

Job Loss. A substantial literature studies the effects of job loss in the context of mass layoffs and establishment closures. Recent research include Lachowska et al. (2020), Schmieder et al.

[3] Almlund et al. (2011) provide an excellent survey of the evidence on the predictive power of personality in the labor market.

(2018), and Huttunen et al. (2011). Several papers have also studied heterogeneous treatment effects among the displaced. For example, von Wachter and Handwerker (2009) and Hoynes et al. (2012) find that job loss is less costly for the college educated. More recently, Kauhanen and Riukula (2019) find that individuals working in occupations with social-intensive tasks before the shock experience the smallest drops in earnings and employment relative to workers in high routine, manual, and cognitive occupations. Our findings complement this result: while Kauhanen and Riukula (2019) compare individuals across occupations, we show that also within occupations, the more extraverted individuals adapt better. The closest papers to our study are Seim (2019) and Dahlberg et al. (2021). Seim (2019) documents that cognitive and non-cognitive skills do not predict faster recovery from job loss using Swedish military-enlistment data. One possibility for the different result could be that our measures capture more precisely the type of skills that help workers adapt; for example, we find that conscientiousness does not predict faster recovery from job loss, while extraversion does. On the other hand, using the same Swedish data but focusing on military personnel affected by military-base closures, Dahlberg et al. (2021) report that non-cognitive skills predict shorter unemployment spells.

2.2 Data

This paper combines several data sources using unique person identifiers.[4]

2.2.1 Psychological Measurement

Data for psychological traits, personality and cognitive skills, are obtained from the Finnish Defence Forces (FDF), which has tested all military conscripts since 1955. The available data cover 79% of Finnish men born 1962–1979 (n = 489,252). These data are the basis for our analysis sample. The FDF data are described in more detail in Appendix B.2.

2.2.1.1 The Data Source

Military conscription in Finland between 1962 and 1979 was universal and granted relatively few exceptions. Finnish men are drafted in the year they turn 18 and most start their service at age 19 or 20. Military service lasts for 6–12 months. Most conscripts do not stay to serve at the military, but continue to civil workforce or studies. FDF uses psychological tests to assess conscripts' suitability for non-commissioned officer training that takes place during the military service.

Both personality and cognitive ability tests are typically taken in the second week of military service in a 2-h paper-and-pencil format in standardized group-administered conditions. The personality test contains 218 statements with a response scale of yes/no. The cognitive test contains 120 multiple-choice questions. The test questionnaires have been unchanged for the timeline of the study, and the scores are designed to be comparable across cohorts.

[4]The data are described in more detail in the Appendix B.2.

2.2.1.2 Test Content

The raw data provide test scores for 8 personality dimensions and 3 cognitive-skill dimensions.

The measured personality traits are: sociability, activity-energy, self-confidence, leadership motivation, achievement motivation, dutifulness, deliberation, and masculinity. The personality test is similar to and based on the *Minnesota Multiphasic Personality Inventory (MMPI)*. The raw scores of the data are a count of yes/no answers that are consistent with the measured trait. For example, a "yes" answer to a statement: "I enjoy spending time with other people", gives a one point toward the sociability score.

The measured cognitive skills are visuospatial, arithmetic, and verbal reasoning. The visuospatial test is similar to Raven's Progressive Matrices (Raven and Court, 1938). The FDF cognitive ability test is similar to the *The Armed Services Vocational Aptitude Battery (ASVAB)*, administered by the *United States Military Entrance Processing Command*. Each correct answer gives a one point toward each cognitive skill measure.

The Appendix provides basic descriptive statistics on the raw personality and cognitive data. Figure B-1 shows the density distributions of each personality and cognitive measure. Both cognitive and personality test scores contain ample variation; for example, there are both people with high and low scores of dutifulness. Table B.1 shows the cross-correlation matrix between the raw personality measures, cognitive skills, education, and prime-age income measures. Personality traits and cognitive scores are strongly correlated within their domains. Correlations across cognitive scores and personality traits are modest.

2.2.1.3 Dimension Reduction

We conduct an exploratory factor analysis to determine a way to reduce dimensionality in our personality and cognitive data. The aim is to isolate distinct personality traits from the relatively high-dimensional data (11 psychological variables). To what extent are the measured personality traits distinct from cognitive skills and each other? The factor-based approach allows us to construct stable variables, avoid multicollinearity between the traits, and reduce measurement error. Based on the analysis described below and guided by evidence from personality psychology, we decide to use a three-factor model, visualized in Figure 2-1. This factorization differentiates between cognitive ability and two personality factors related to extraversion and conscientiousness (*interpersonal* vs. *intrapersonal* traits).

The eigenvalue plot from our exploratory factor analysis is provided in Figure B-2. The eigenvalue plot supports the idea of dimension reduction: our raw data have 11 dimensions but 5 factors are enough to account for almost all of the variation. The eigenvalues suggest that we should retain at most 5 factors. Our decision to use only three factors is based on the objective to reduce the dimensionality of the data while still retaining interpretability. With three factors, the 11 traits divide quite cleanly into cognitive ability plus two out of the widely used "Big Five" personality traits.[5]

In psychology research, the Big Five traits are often further divided into subtraits (facets) that are measured with standard questionnaires (Corr and Matthews, 2020). While our data do not come from such a standard questionnaire, most traits in our data correspond to a subtrait of one or more Big Five traits. Sociability, activity, confidence, and leadership are subtraits associated with extraversion. Deliberation, dutifulness, and achievement motivation are subtraits associated with conscientiousness. Masculinity is not associated with Big Five traits in any common operationalization of the Five Factors Model.

The factor loadings from the common factor analysis are reported in Table B.2. We use an oblique rotation where the factors are allowed to be correlated. In a two-factor model, the cognitive and personality test scores load on distinct factors, as shown in Jokela et al. 2017. In a three-factor model, the extraversion-related scores (sociability, activity, confidence, leadership) load onto a separate factor and the conscientiousness-related scores (dutifulness and deliberation) load onto a separate factor. The remaining two raw measures do not load strongly onto either factor: Achievement aim loads onto the extraversion-related factor (despite being associated with conscientiousness) but has the lowest loading within that factor and, at the same time, the third highest loading on the conscientiousness-related factor. Our interpretation is that the FDF achievement aim measure combines both external and internal motivations for achievement. Masculinity has a low factor loading in any of the factors and a high uniqueness score.[6]

Based on the close grouping of the subtraits (in terms of factor loadings) with their corresponding Big Five domains, we proceed to refer to the two personality factors as extraversion and conscientiousness. Because our measures do not correspond perfectly with any particular operationalization or a survey of the Big Five traits, this terminology is not exact. However, Jokela et al. 2017 show that using a separate survey to capture the Big Five traits in convenience sample, the FDF measures are correlated with extraversion and conscientiousness in the expected directions.

For the main analysis, we construct variables from the three-factor model by estimating the factor scores for each individual and normalizing the variables to have zero mean and unit standard deviation.

2.2.2 Labor Market, Education and Demographics

The paper takes advantage of the detailed longitudinal register data on the full Finnish population of individuals and firms compiled by the *Statistics Finland* from multiple sources. Plant, firm, industry, local-level, and similar measures are computed from the full data, containing all persons in Finland. We manually harmonize all occupation, education, industry, and geographical classifications to be consistent over time.

[5] These traits are extraversion, conscientiousness, neuroticism, openness to experience, and agreeableness.

[6] Allowing for four factors essentially adds an extra factor for masculinity. To keep the analysis tractable, we do not include masculinity as a separate factor in our analysis. In a separate paper (Izadi and Tuhkuri, 2021b), we analyze masculinity in a more detail.

The register data provide information on demographics, labor market status, earnings, occupation, industry, firm and establishment identifiers, and county of residence and birth, for all Finnish residents 1987–2019.

Income data are obtained from the *Finnish Tax Authority.* The primary earnings measure is the yearly labor earnings from the primary employment relationship. We measure 'prime-age' earnings as the average annual labor-market earnings during ages 35–38. We deflate all values to 2010 Euros using the Statistics Finland CPI and drop the observations with zero prime-age earnings from the earnings analyses (less than 1%).

The Register of Completed Education and Degrees contains exact information on the educational degrees the individual has obtained, including both the level and field, and the date at which the degree was granted. All degrees completed in Finland are generally recorded in these data. When we use education just as a control variable, we include only education level and field fixed effects. Otherwise, we map degrees to years of education according to their official length (e.g., a master's degree equals 17 years of education). GPA at the 9th grade is measured from the *Secondary Education Application Register* and high-school graders from the *Finnish Matriculation Examination Board Register.*

2.3 Descriptive Evidence

We begin the analysis by relating three psychological factors (cognitive ability, extraversion, and conscientiousness) and education to labor-market outcomes in the cross-section and demonstrate their relationships to each other.

This section shows that cognitive ability and education are important in predicting labor market success relative to extraversion and that conscientiousness has significant predictive power in the labor market. We later contrast this finding by showing the opposite order of importance in response to a labor market shock, where extraversion becomes the best predictor of adaptation.

In our measurement, we draw a distinction between interpersonal vs. intrapersonal traits. The factor variable extraversion measures traits that affect relationships between people. The factor variable conscientiousness measures traits that work primarily within the person. We also make a distinction between a person's type vs. skill. The main difference is that type is a set of attributes fixed at the point of measurement, while skill is endogenous to the type. We view personality traits and cognitive ability as a type and education as a skill. Due to this endogeneity, we focus on regressions where education is excluded, but for a reference, also provide estimates where it is included.

2.3.1 Cross-Correlations

Table 2.1 presents the cross-correlations between the main factor variables, prime-age earnings, and the 9th grade GPA. The main observations are: (1) cognitive ability, education, and school GPA are relatively closely correlated with each other ($\rho > .5$), (2) extraversion and conscientiousness have relatively low correlations with each other and with cognitive ability, education, and GPA ($\rho < .35$), and (3) all traits positively correlate with earnings.

2.3.2 Cross-Sectional Evidence on Earnings

Table 2.2 presents the standard cross-sectional estimates of the predictive labor-market returns to each trait. The cross-sectional estimates are from specification:

$$Y_i = \beta \times \text{Trait}_i + \gamma_i + \varepsilon_i. \tag{2.1}$$

The outcome is log prime-age earnings, and Trait is a vector of traits.[7] The model controls for birth-year fixed effects (γ_i). We present three versions: (1) the estimates for each factor variable separately, (2) with all factor variables, and (3) with all factor variables and education. The first four columns reveal in regression form the same cross-correlation pattern as in Table 2.1. One SD increase in extraversion or conscientiousness is associated with about a 20% increase in prime-age earnings. The same increase in cognitive ability is associated with a 35% increase in earnings. Column 5 shows that once all three are included in the same regression, coefficients for extraversion and conscientiousness are halved, but cognitive ability decreases little. When the years of education are added in Column 6, the coefficient for conscientiousness and cognitive ability decrease, but extraversion remains unchanged relative to Column 5. The connection between personality traits, education, and earnings in the cross section is analyzed in Izadi and Tuhkuri (2021b).[8]

Figure 2-2 visualizes the conditional expectation function (CEF) for each factor. The outcome is prime-age earnings. The visualization of the CEF groups the x-axis variable into equal-sized bins, computes the mean earnings within each bin, and creates a scatterplot of these data points. The visual evidence confirms that cognitive ability, extraversion, conscientiousness, and education all positively predict prime-age earnings.

2.3.3 Cross-Sectional Evidence on Adaptive Behaviors

Table 2.3 presents the cross-sectional estimates focusing on a wider set of outcomes that measure potentially adaptive behavior in the labor market. The main set of outcomes measure switching of occupation, industry, firm, and educational status. We operationalize these measures as the total count of switches between ages 28–38. We also provide an estimate for employment over time, operationalized as a yearly indicator for being employed over ages 28–38.[9] To preserve space, we use a single specification that estimates the heterogeneous returns of all factors jointly.

As shown, individuals are employed on average 10 years out of the 11 year period and switch occupation, industry, and establishment .5–1 times. The results show that conscientiousness is

[7] We measure 'prime-age' earnings as the average annual labor-market earnings during ages 35–38. We deflate all values to 2010 Euros using the Statistics Finland CPI and drop the observations with zero prime-age earnings from the earnings analyzes (less than 1%).

[8] We find that specific traits are negatively associated with education but positively with earnings.

[9] Note that while we observe all our sample persons at the prime age, we do not observe all persons between ages 28–50: our labor-market data are available between 1987–2018 and the sample covers birth cohorts 1962–1979.

positively associated with employment but negatively with switching: conscientious individuals find a job and stick to it. This is contrasted by extraversion, which predicts frequent switching but not particularly high employment. Lastly, high cognitive ability predicts both high cumulative employment and frequent switching of occupation and firm. We will return to these outcomes in the last section to show that the patterns are different during times of economic distress.

2.4 Mass-Layoff Evidence

This section analyzes how different dimensions of human capital—personality traits, cognitive ability, and education—mediate how individuals adapt to a negative labor-demand shock at their firm. The firm-level shock is a case study that compares *stable* versus *unstable* times for an individual in the labor market. We look at both short and long-term adaptation.

To define and measure a negative firm-level shock, we focus on a mass-layoff event. Mass layoff is an episode where a firm or an establishment simultaneously lays off a large share of its workers (see, for example, Jacobson et al., 1993). We analyze the reduced-form effects of a firm-level shock, and by doing so, depart from the standard focus on (endogenous) job loss. The main reason is that our focus is on adaptation; selection into exit from the plant is in principle an essential part of the mechanism. The unit of observation for measuring the mass layoff is the establishment; for simplicity we refer to it as the firm.

Our main analysis explores how the returns to different dimensions of personality and skills depend on whether or not the person was subject to the event. This is a heterogeneity-based approach for analyzing how the effects of a mass-layoff event depend on the characteristics of the individuals exposed to the event.

We define a treatment group as workers who experienced a mass layoff shock and had a strong attachment to the labor market before the shock. We construct a counterfactual by matching workers who experienced a mass layoff in a given year to a comparison group of workers who were similar based on a rich set of characteristics but did not experience the shock. We compare these matched workers—the treatment and the control group—using an event study type specification. The event study shows whether the two groups followed similar trends leading up to the event and identifies how their outcomes diverged after the mass layoff.

2.4.1 Setup

2.4.1.1 The Mass Layoff Event

We define the mass-layoff event by using the following criterion: The plant reduces its employment by at least 30% between year t and $t+1$. This definition includes full closures. To reduce measurement error, we require that no more than 50% of the exiting employees continue in the same new plant after the event (we exclude "false events"). For full closures, we require that the firm does not re-appear in the data. We use the term "mass layoff" to refer to both mass layoffs where the plant continues its operations and full plant closures.

2.4.1.2 The Treatment Group

The basis for the treatment group is the workers exposed to a mass layoff. We define them as a set of workers that were working at a plant j in year $t-1$ when the plant had a mass layoff between year t and $t+1$. The timing is defined this way to ensure that the sample of workers remaining in the firm before the mass layoff is not excessively selected.[10] In the figures, we we label $t-1$ as period zero, and thus the event happens between periods one and two.

The pool of potential treatment units is Finnish men born between 1962 and 1979 that have military test records available. We consider event years 1995–2010.[11] This period includes all phases of the business cycle. Macroeconomic conditions are shown to have a large influence on treatment effect estimates in mass-layoff settings (Davis and von Wachter, 2011; Schmieder et al., 2018). We do not focus on business cycle variation, but our estimates can be viewed as long-term averages concerning the state of the economy.

To construct the treatment group, we apply a set of sample restrictions. The idea is to focus on workers that had a strong attachment to the labor market and a stable employment relationship before the shock. These are workers that switch from a stable to an unstable labor-market situation. To capture this idea, we focus on prime-age workers and require that the worker is at least 35 years old in the year before the mass layoff $t-1$, has been continuously employed from $t-6$, and continuously employed at the given firm from $t-4$.

We restrict the sample to establishments with 5–2000 workers. For the mass-layoff events that are not full closures, we require that the plant had at least 20 workers in year $t-1$. We apply a floor to the plant size because the concept of a mass layoff or plant closure requires at least a few workers, for which the event was relatively unanticipated. Micro establishments are also excluded since we aim to focus on workers that are paid employees rather than entrepreneurs or family members. We apply a limit to the plant size because plants with over 2000 employees tend to be outliers or multi-plant firms classified as single plants.

To restrict the influence of outlier observations, we exclude top and bottom 1% of labor-income earners from the final sample and observations where the earnings are more than 3 times higher than the base year earnings. We apply no industry or firm-type restrictions. We focus on the first mass layoff for each individual that satisfies the data restrictions, and require no previous mass-layoff events between $t-5$ and $t-1$.

2.4.1.3 The Matched Control Group

To construct a counterfactual for the treatment group, we use coarsened exact matching (CEM). The pool of potential control units is all male workers with military records but with no mass layoff event in a window from $t-5$ to $t+8$, the estimation window. We use the event time $t-1$ to measure the match variables.

[10] There is a trade-off: The closer we move to the event, the stronger the workers' attachment to the firm. The further we move from the event, the less likely the workers will have anticipated the event.

[11] Before 1995 our first cohort would be too young, and after 2010 our post-period would be too short.

We perform the match in three steps: (1) We apply the treatment-group restrictions to the pool of potential control units. (2) We match on exact characteristics: year, age, tenure, industry, and firm size.[12][13] (3) We perform a caliper match based on pre-period earnings to select the closest matches within the set of exact matches. In the case of a tie, we choose the control person with a non-missing occupational code. We set the ratio of treatment to control units to 1:5.[14] The match is performed with replacement.

Focusing on the matched control group that never receives treatment reduces the problems arising in estimating dynamic treatment effects when the comparison group consists of units treated at different points in time (Sun and Abraham, 2021; Goodman-Bacon, 2021).

2.4.1.4 Descriptive Statistics

Table B.3 presents worker-level descriptive statistics, and Figure B-3 compares the distributions of main outcomes for the treatment and the control groups in the first pre-period. The treatment group has 18,005 individuals, the control group has 89,360 individuals. The treatment and control groups are similar on a wide set of outcomes, although similarity in levels is not required in later analysis.

Table B.4 collects plant-level information. The sample contains 3,639 treatment plants, and 31% of the events are full closures. The treatment firms' typical employment reduction is 49%, while the control group firms typically increase employment by 3.7%. The typical industries in the sample are manufacturing of electronics, machine, paper, and wood; construction; wholesale trade; and transportation. The typical occupations are machine operators; metal, machinery and related trades workers; construction and related workers; science and engineering professionals and associate professionals; and drivers and mobile-plant operators.

2.4.2 Estimates

2.4.2.1 Mass Layoffs' Effects on All Workers

This section provides the baseline estimates for the effects of the mass layoff event on workers' labor-market performance. We use three tools: raw means, event-study estimation, and pooled difference-in-differences estimates.

The design is visualized in Figure B-4, which plots the raw means of employment and earnings for the treatment and the matched control group over the event time. The treatment group experiences a sharp decline in both outcomes right after the event. The control group displays mean reversion when sample restrictions are lifted after the event. The figure underscores that

[12] Coarsened classes: year in years, age in 2-year bins, tenure in years until 7 and then 8-10, 11-20, 20-, industry in harmonized sectors (7), firm size in 0–25, 26–50, 51–100, 101–250, 251–500, 501–1000, 1001–2000.

[13] The match on tenure is important: To be subject to a mass layoff or plant closure, the worker needs to be employed. The longer the worker is employed in a given firm, the higher the likelihood of being subject to a mass layoff or plant closure event. Compared to the full population, those subject to a mass layoff or plant closure are positively selected in terms of employment history and income.

[14] 99% of the treatment units have 5 matched control units that fulfill the criteria.

being continuously employed is unlikely to be the correct counterfactual for the treatment group (Krolikowski 2018).

To quantify the differences between the treatment and control groups, we estimate the following event-study specification:

$$Y_{ijt} = \alpha_{iy} + \gamma_t + \sum_{t=-5}^{8} \delta_t \times \text{Treat}_i + \boldsymbol{X_{ijt}\theta} + \varepsilon_{ijt}. \tag{2.2}$$

The main outcomes Y_{ijt} are earnings (relative to the base year) and employment (in general and in the baseline firm j). The index t denotes the event-time, i the individual, j the establishment, and y the event year. The specification includes fixed effects for the individual $\times$ event year (α_{iy}), and time relative to event (γ_t). The term $\boldsymbol{X_{ijt}}$ denotes potential other time varying controls such as age. To account for unobserved common shocks, we cluster standard errors at the establishment level. We omit event time $t-1$ as the reference category. The key identifying assumption is the parallel trends of potential outcomes. Conditional on parallel trends of potential outcomes, the $\delta_{\mathbf{t}}$ estimate the causal effects of the shock on earnings and employment at a given time.

Figure 2-3 reports the $\delta_{\mathbf{t}}$ estimates. Pre-trends are absent in the figure (by construction of the matched control group in the case of employment). Immediately after the event, workers' earnings decrease by 10% on average relative to the event year. The decrease persists for at least the following eight years. Employment also decreases by 9% among the affected but regains about half of that loss during the first five years after the event.

To combine the event-study coefficients into a single treatment effect estimate, we also estimate a pooled difference-in-differences specification:

$$Y_{ijt} = \alpha_{iy} + \delta_t \left(\text{Treat}_i \times \text{Post}_t\right) + \gamma \text{Post}_t + \boldsymbol{X_{ijt}\theta} + \varepsilon_{ij}, \tag{2.3}$$

where $\text{Post}_t = 0$ before the shock ($t \in [-5, 0]$) and $\text{Post}_t = 1$ after the shock ($t \in [2, 8]$). Treat_i main effect is absorbed by the individual $\times$ event year (α_{iy}) fixed effects. We exclude the first period from these estimations because treatment is defined at period zero, whereas the actual event happens between periods one and two. The results for earnings and employment are reported in Table 2.4. On average, earnings fall by 9.8% in the post-period relative to the event year as a consequence of the event. Employment falls on average by 6.2%.

2.4.2.2 Mass Layoffs' Effects Depending on Workers' Characteristics

This section estimates the heterogeneous effects of different psychological traits on workers' labor-market performance, conditional on whether the workers were exposed to the mass layoff event. To approach this goal, we use three tools: raw quantile means, heterogeneous effects in an event-study framework, and pooled difference-in-differences. The main outcomes Y_{ijt} are earnings (relative to the base year) and employment (in general). We focus on the earnings relative to the baseline since it (1) captures the idea of adaptation and recovery, (2) allows to use zero-values, and (3) is intuitive to interpret in percentages.

To set the stage, we present the raw means in the top and bottom quartile (top vs. bottom 25% within the mass-layoff sample) of each trait separately for the treatment and control groups. Figure B-5 visualizes the results for the main outcomes: earnings and employment. The figure shows that the immediate employment drop for each trait is smaller for the top quartile individuals than for the bottom quartile individuals. College-educated individuals also suffer a much smaller employment drop than non-college-educated individuals with a comparable magnitude. The differences between the top and bottom groups are less clear in earnings due to pre-treatment level differences between the groups. The raw means also do not consider the partial correlations of the factor variables between each other. To address these issues and estimate the magnitudes of these differences, we next estimate the differential effects of the shock in an event-study framework.

We augment Equation 2.2 by adding a triple-difference interaction term for each trait:

$$Y_{ijt} = \alpha_{iy} + \gamma_t + \sum_k \sum_{t=-5}^{8} \delta_{tk} \times \text{Treat}_i \times \text{Trait}_{ik} + \boldsymbol{X_{ijt}\theta} + \varepsilon_{ijt}. \tag{2.4}$$

The index t denotes the relative event-time, i the individual, j the firm, y the event year. All lower-order (pairwise) interactions are included in $\boldsymbol{X_{ijt}}$. To account for the residual correlation between the factors, we estimate each of the three traits—cognitive ability, extraversion and conscientiousness (indexed by k)—jointly in the same regression. Education is estimated in a separate regression without including traits. We estimate traits separately from education because education is potentially influenced directly by traits as shown in Izadi and Tuhkuri (2021b). To account for unobserved common shocks, we cluster standard errors at the establishment level.

Figure 2-4 presents the results for earnings and employment. Each line shows the $\delta_\mathbf{t}$ estimates for the corresponding trait. For example, the green line in the first panel of Figure 2-4 shows that in period three, extraverted individuals (one standard deviation above the sample mean) have about 2 percentage points smaller earnings losses than individuals with average traits. In other words, the negative effect of the mass layoff on earnings is about two percentage points smaller for extraverted individuals, holding cognitive ability and conscientiousness fixed. Compared to the baseline of 10%, this amounts to about a 20% reduction in the effect per standard deviation of extraversion. For extraversion, this reduction extends to the end of the observation period. In contrast, while individuals with high cognitive ability also experience a smaller initial hit on earnings, they are caught up by the average individual by period eight. Finally, conditional on extraversion and cognitive ability, conscientiousness does not predict adaptation to the shock. In Figure 2-5, education behaves similarly to cognitive ability. It has a transitory moderating influence on the magnitude of the earnings reduction, which then fades away in later periods. An additional year of education is worth about one standard deviation of cognitive ability in terms of reducing the short-term effect of mass layoff.

The right panels in Figures 2-4 and 2-5 present the $\delta_\mathbf{t}$ coefficients for employment as the outcome. The results are similar to earnings. Extraverted individuals experience a permanently smaller drop (up to two percentage points) in employment after the shock relative to the average individual, whereas high cognitive ability and education predict more transitory reductions in the negative effects of the shock on employment. Conscientiousness remains a weak predictor of adaptation conditional on other traits. These results should be compared to the baseline estimate of the impact of the shock on employment, which is initially about 9 percentage points. Taken together, the evidence so far suggests that the better adaptation to the unexpected mass-layoff shock, which is enjoyed by the extraverted, and to a lesser extent, the highly educated and those with high cognitive ability, is associated with the employment margin.

The quantity of interest can be viewed as a triple differences estimate, where the third difference comes from the variation in traits. We estimate the following specification, which provides a single estimate for the trait-dependent differences in response to the shock:

$$Y_{ijt} = \alpha_{iy} + \sum_{k} \beta_k \left(\text{Trait}_{ik} \times \text{Treat}_i \times \text{Post}_t\right) + \gamma \text{Post}_t + \boldsymbol{X_{ijt}\theta} + \varepsilon_{ijt} \tag{2.5}$$

where $\boldsymbol{X_{ijt}\theta}$ further includes a full set of interaction terms between the Trait, Treat, and Post indicators. The Trait_{ik} and Treat_i main effects are absorbed by the individual $\times$ event year fixed effects (α_{iy}). The triple-interaction terms correspond to a weighted average of the post-event estimates in the previous figures. Table 2.5 presents the results for earnings. The first two columns correspond to the specification used in Figure 2-4, where traits are estimated jointly, but education is estimated separately. The coefficient for extraversion is 2%, as noted earlier. The coefficients for cognitive ability and education are lower than in the first post-periods due to their declining effect. Column 3 estimates education jointly with the psychological traits. In this specification, the coefficients for cognitive ability and education have decreased relative to Columns 1 and 2, indicating that they partly capture the same heterogeneity. Including education does not change the coefficient for extraversion.

An important caveat in this analysis is the causal interpretation of the coefficients in Equation 2.4. Briefly, they do not have one. The arguably exogenous variation in our setting comes from the unexpected mass layoffs in firms. That gives the baseline estimates in Section 2.4.2.1 a causal interpretation. However, without additional assumptions, the coefficient of interest in Equation 2.4 is strictly descriptive. In particular, personality traits, cognitive ability, and education can influence the individual's response to the shock indirectly through selection on unobservables, such as occupational choice and selective layoffs. Maybe extraverted individuals work in occupations or industries with less competitive labor markets where re-employment is easier? Column 4 in Table 2.5 includes controls for occupation and industry in period 0. The categorical dummies are fully interacted with Treat and Post to allow the treatment effect to vary across occupations and industries. Including these fixed effects slightly reduces the coefficient of extraversion, indicating that a small part of the positive effect of extraversion may be driven by pre-treatment selection into occupations and industries.

Table 2.6 displays the estimation results for employment, which closely follow the earnings estimates. We take this as suggestive evidence that the heterogeneous effects that we find for psychological traits are primarily mediated by employment opportunities instead of changes in wages. In the next section, we look at different behaviors which could explain the heterogeneous effects.

2.4.2.3 Mass Layoffs' Effects on Outcomes Related to Adaptive Behaviors

This section analyses the potential mechanisms that lead to different adaptive responses between different kinds of individuals. To explore the potential mechanisms of adaptation—the channels through which different psychological traits influence recovery and resilience—we look into an extended set of outcomes. How do people change their labor-market behavior after an unexpected labor-market shock? Why do extraverted individuals experience smaller drops in earnings and employment?

We start by estimating the baseline Equation 2.2 for four new outcomes: plant exit, occupation change, industry change, and re-education. Figure 2-6 presents the results. The first panel shows the event-study coefficients for plant exit probability, or the "first stage," of our baseline event study. Individuals employed in the plant before the mass layoff are 50 percentage points more likely to exit their plant in period two than the control group. However, as noted earlier, the dis-employment effect of the event is only 9% in the short term. The vast majority of laid-off individuals find re-employment during the same year somewhere—most individuals adapt to the shock by finding new employment soon after.

The second panel shows the probability of changing occupations. Change is measured relative to period zero. The treatment group has consistently about 9 percentage points higher rate of occupational change relative to the treatment group. As a benchmark, the occupational change rate in period two in the control group is about 23%. This shows that occupational change is an important adaptive margin. However, industry change is even more typical. The effect of a mass layoff on the probability of changing industry is almost 25 percentage points against a baseline of 7% for the control group in period two. An important caveat is that the resolution of the occupation and industry categories influences the baseline magnitudes: We have 45 occupation categories and 136 industry categories in our sample. The final panel shows the effect of the shock on the probability of re-education. We determine re-education as obtaining a new degree that is either from a different field or more advanced than the individuals' current degree. Over the long term, the effect of the shock on the re-education rate is 2 percentage points. The baseline re-education rate in the control group in the last period is 5.5%.

Overall, we have identified four potentially important margins of adaptation: job retentions at the original establishment, industry change, occupation change, and re-education. Next, we will analyze how different traits and education levels interact with these margins, and estimate Equations 2.4 and 2.5 for this new set of outcomes.

We first focus on plant exit. The first panel in Figure 2-7 and Column 4 in Table 2.7 show the estimates. The green line shows that even in the long term, extraverted individuals are less likely to exit their original plant than the average individual in our sample. However, the magnitude is relatively small: while the baseline exit probability is 50 percentage points higher in the treatment group, extraversion reduces this at most by two percentage points.

Why are the extraverted individuals more likely to survive the mass layoff and keep their job at the firm? One possibility is that they are working in different occupations at the firm. Table 2.8 controls for occupation, education, and industry in the pooled triple-difference specification. Column 4 shows that the differential plant exit rate decreases to less than one percentage point with the controls included. At least half of the job retention advantage among the extraverted appears to be explained by selection.

The story for cognitive ability is different (the blue line). High cognitive-ability individuals seem to anticipate the layoff and are *more* likely to exit the plant before the layoff (mass layoff happens between periods one and two). However, the estimate is small in magnitude and not statistically significant. After the first period, there is no significant difference between the exit rates of high cognitive ability individuals and the average.

Now we look into industry and occupation changes. The second and third panels in Figure 2-7 show that industry and occupation changes induced by the shock are about 2 percentage points *less* common among the extraverted. That is, surprisingly, extraversion does not predict more frequent re-allocation. Recall that in the cross-sectional estimates presented in Section 2.3 (Table 2.3) we found that extraverted individuals are more likely to work in multiple firms, occupations, and industries during their careers. But the shock disproportionately induces the extraverted individuals to adapt by seeking employment in the same type of occupations and industries as before the layoff. This effect is partly also expected as they retain their job at the firm but the occupation result is still robust to controlling for baseline occupation and industry in Table 2.7. The patterns for cognitive ability and education (Figure 2-8) are similar in terms of industry and occupation changes.

For re-education, both predictors of positive adaptation, extraversion and cognitive ability, predict lower re-education rates. Some of the effects may be driven by having less room for educational upgrading because of higher baseline education rates among these individuals, and the effects are marginally significant.

In summary, the traits that predict adaptation—especially extraversion—seem to help workers find re-employment faster in a similar occupation and industry they worked in before. This result is not entirely driven by higher job retention or selection into specific pre-shock careers. Faster adaptation is associated with lower re-allocation in terms of industry, occupation, and education.

2.5 Conclusion

Labor markets are in constant change. These changes put people in situations that require resilience and adaptation. This paper analyzes how individuals' resilience to a labor-market shock varies by their psychological profiles. We use mass layoffs at their workplaces as a case study. We use standardized personality test results from the Finnish military conscription to construct measures of cognitive ability, extraversion, and conscientiousness. We find that extraversion is a powerful predictor of recovery. Even in the long term, extraverts experience significantly smaller adverse effects from this shock. Our results are driven by faster re-employment rather than wage growth or changing industry and occupation after the shock. Extraverts are slightly more likely to retain their employment at a mass-layoff establishment, but that is not the primary driver of our result.

Classic theoretical research in economics (Nelson and Phelps, 1966; Welch, 1970; Schultz, 1975) emphasizes the role of human capital as the capacity to adapt, in contrast to its productive value at work. We contrast the adaptive vs. productive value by comparing the value of personality traits and skills in the cross-section vs. labor-market shock. In the cross-section, cognitive ability is the best predictor of earnings, while conscientiousness and extraversion are approximately equally important. In contrast, in a mass-layoff situation, extraversion is the best predictor of recovery. Cognitive ability is still important, but conscientiousness does not predict better adaptation. These observations demonstrate that the characteristics that predict adaptation are different from those that predict labor-market success overall. The paper also contributes to the long-standing debate on person vs. situation as determinants of individual behavior (see, for example, Ross and Nisbett, 1991): Person and situation together matter when estimating the economic benefits of individual traits.

Recent research in economics analyzes the value of social skills in the labor market (Deming, 2017). We find that the value of extraversion appears to be pronounced in situations that require resilience and adaptation. This finding provides a new complementary interpretation for the previously observed economic value of social skills in the labor market (Deming, 2017).

Identifying predictors of adaptation is a first step toward understanding the behaviors and personal characteristics that make people resilient in the labor market. We showed that some salient labor market behaviors, such as pre-shock career choices, are not the likely drivers. Likewise, we showed that extraverts do not markedly differ in post-shock behaviors, such as changing occupations and industries. Further identifying the behaviors that help the extraverts gain re-employment and maintain higher earnings is a natural next step for future research. Recovering from a shock can be related to many skills that are more prevalent among the extraverted. For example, navigating job search and using personal and professional networks in employment search may be easier for extraverted persons.

To the extent that adaptation and resilience are individual skills that can be learned or altered, the findings of this paper could inform policies and research that target the learning of those skills.

Main Figures and Tables

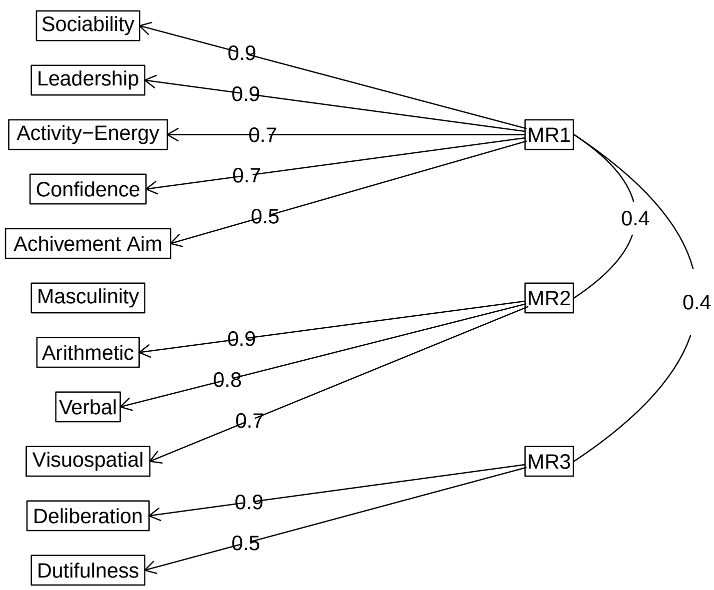

Figure 2-1: Factor Loadings.

Notes: Results from an exploratory factor analysis using three factors with oblique rotation. The numbers on the left indicate the correlation of the test item with the latent factor. The numbers on the right show the correlations between factors. For each test item, only the highest factor loading is shown. MR1 (MinRes solution) is labeled Extraversion, MR2 is labeled Cognitive Ability, and MR3 is labeled Conscientiousness.

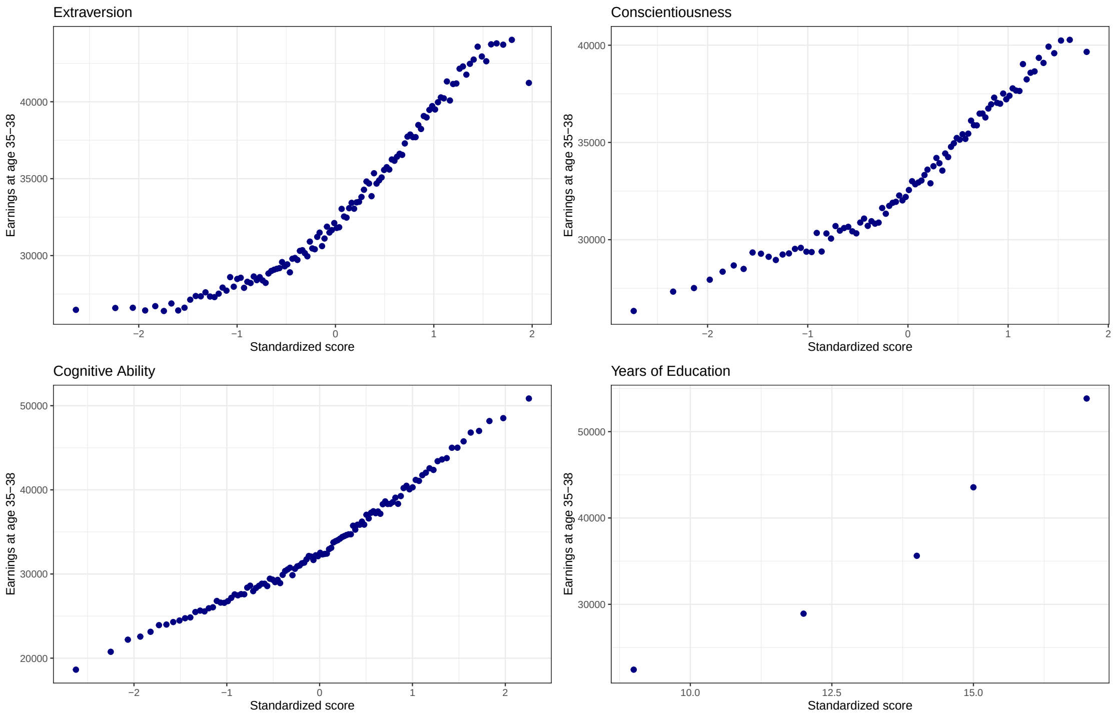

Figure 2-2: Conditional Expectation Functions.

Notes: For the psychological measures, the x-axis is divided in equal-sized bins. Each point represents the mean earnings in that bin in 2010 euros. Earnings are calculated as the sum of labor, and entrepreneurial income averaged over age 35-38. The years of education are computed from the degrees' official lengths (e.g., a high-school degree is 12 years).

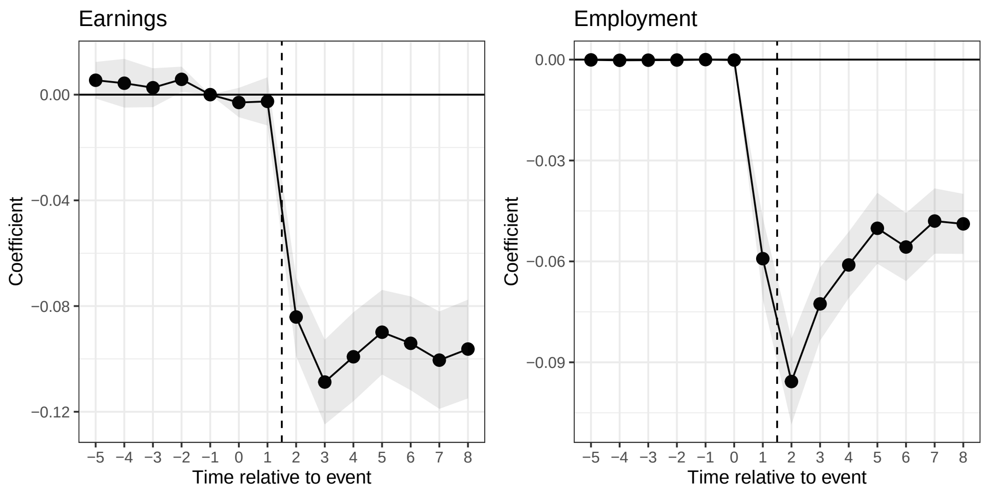

Figure 2-3: Baseline Event-Study Estimates.

Notes: The figure shows the δ_t coefficients from the baseline event-study specification in Equation 2.2. The treatment group consists of workers whose firms experience a mass layoff or closure in period 1. The control group is constructed by matching to workers in firms that do not experience mass layoffs before period 1. Earnings are measured by dividing total labor and entrepreneurial income with period 0 earnings. Employment is binary and takes the value of 1 if the individual is employed during the last week of the year.

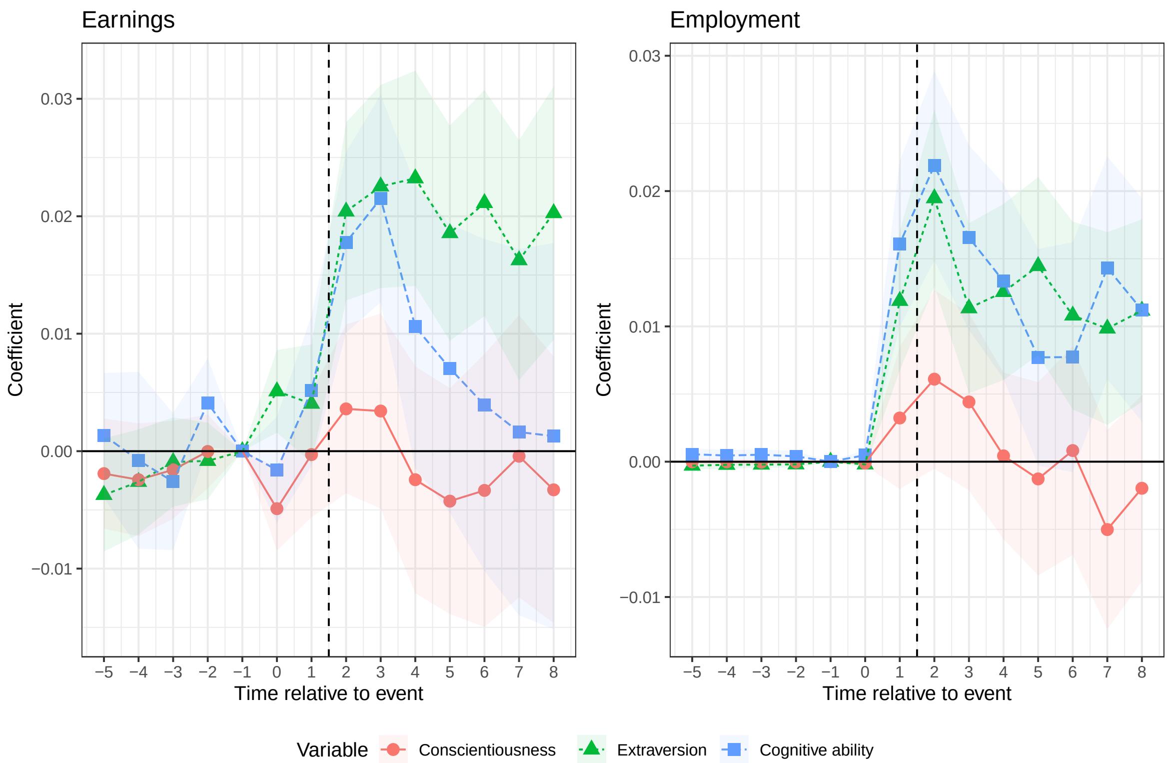

Figure 2-4: Heterogeneous Responses by Trait.

Notes: Each point is a δ_{tk} coefficient from Regression 2.4 for the indicated factor variable. All three factor variables are estimated jointly in the same regression. The left panel is estimated using earnings as the outcome. Earnings are measured by dividing total labor and entrepreneurial income with period 0 earnings. The right panel uses employment as the outcome. Employment is binary and takes the value of 1 if the individual is employed during the last week of the year.

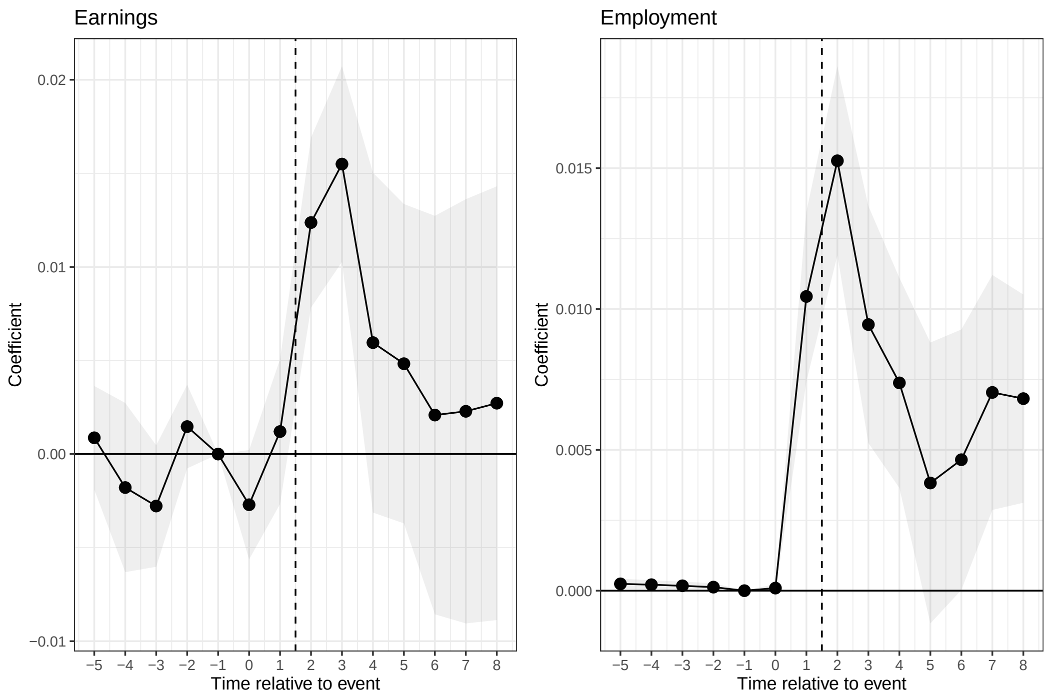

Figure 2-5: Heterogeneous Responses by Education.

Notes: Each point is a δ_t coefficient from Regression 2.4 where Years of Education is used in place of Trait_i. Years of Education is constructed by mapping degrees to their official length (e.g., a master's degree equals 17 years of education). The model is estimated without any of the factor variables. The left panel is estimated using earnings as the outcome. Earnings are measured by dividing total labor and entrepreneurial income with period 0 earnings. The right panel uses employment as the outcome. Employment is binary and takes the value of 1 if the individual is employed during the last week of the year.

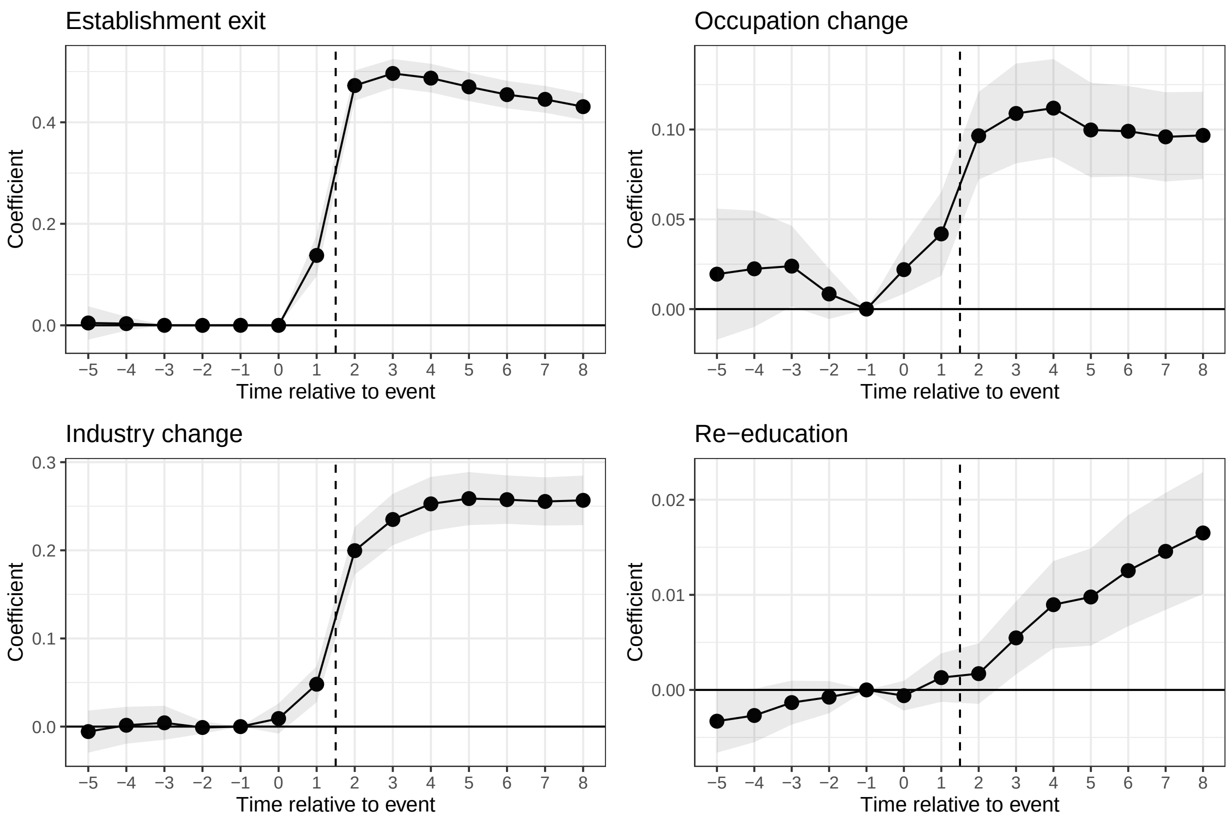

Figure 2-6: Baseline Event-Study Estimates: Adaptive Behaviors.

Notes: The figure shows the δ_t coefficients from the baseline event-study specification in Equation 2.2. The outcome used in the estimation is indicated in the panel name. All outcomes are binary and measured relative to their period 0 value. Re-education takes the value of 1 if the degree does not match the period 0 degree. Industry and occupation are measured only for the employed, which restricts the estimation sample to those employed in the post-period. The treatment group consists of workers whose firms experience a mass layoff or closure in period 1. The control group is constructed by matching to workers in firms that do not experience mass layoffs before period 1.

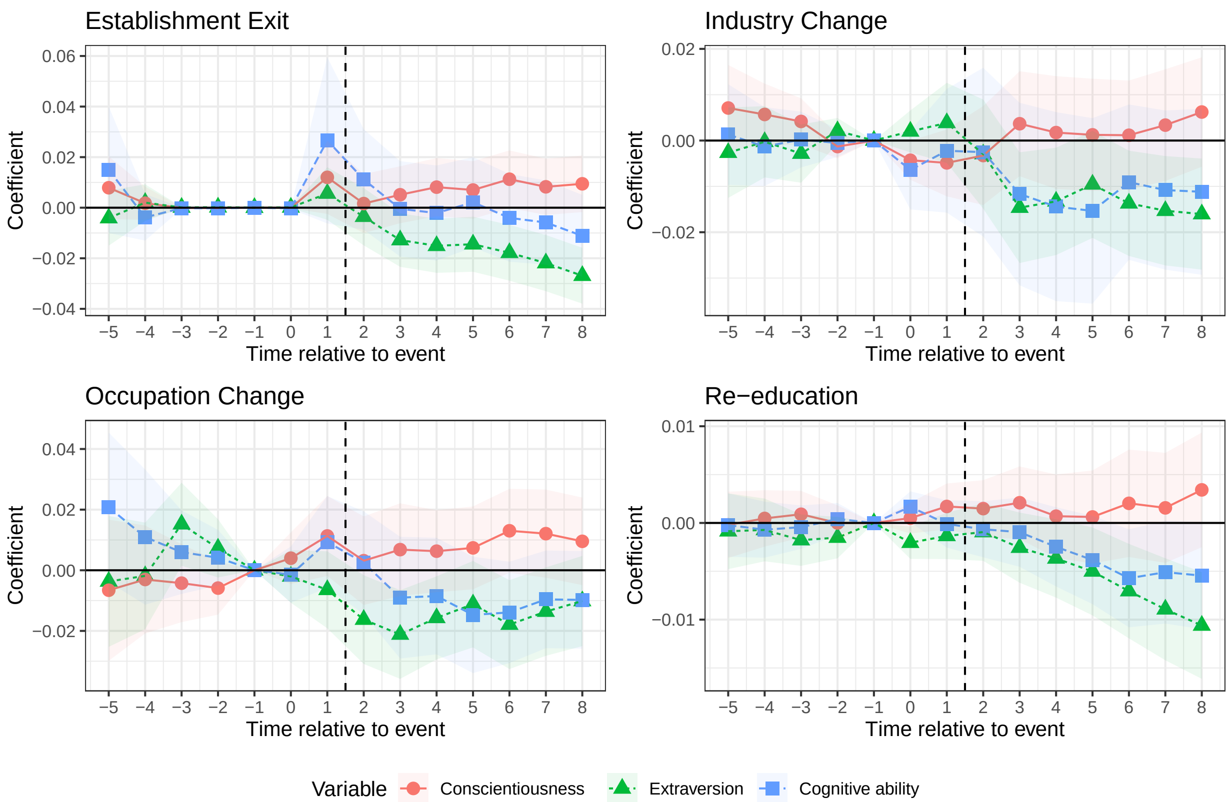

Figure 2-7: Heterogeneous Responses by Trait: Adaptive Behaviors.

Notes: Each point is a δ_{tk} coefficient from Regression 2.4 for the indicated factor variable. All three factor variables are estimated jointly in the same regression. The outcome used in the estimation is indicated in the panel name. All outcomes are binary and measured relative to their period 0 value. Re-education takes the value of 1 if the degree does not match the period 0 degree. Industry and occupation are measured only for the employed, which restricts the estimation sample to those employed in the post-period.

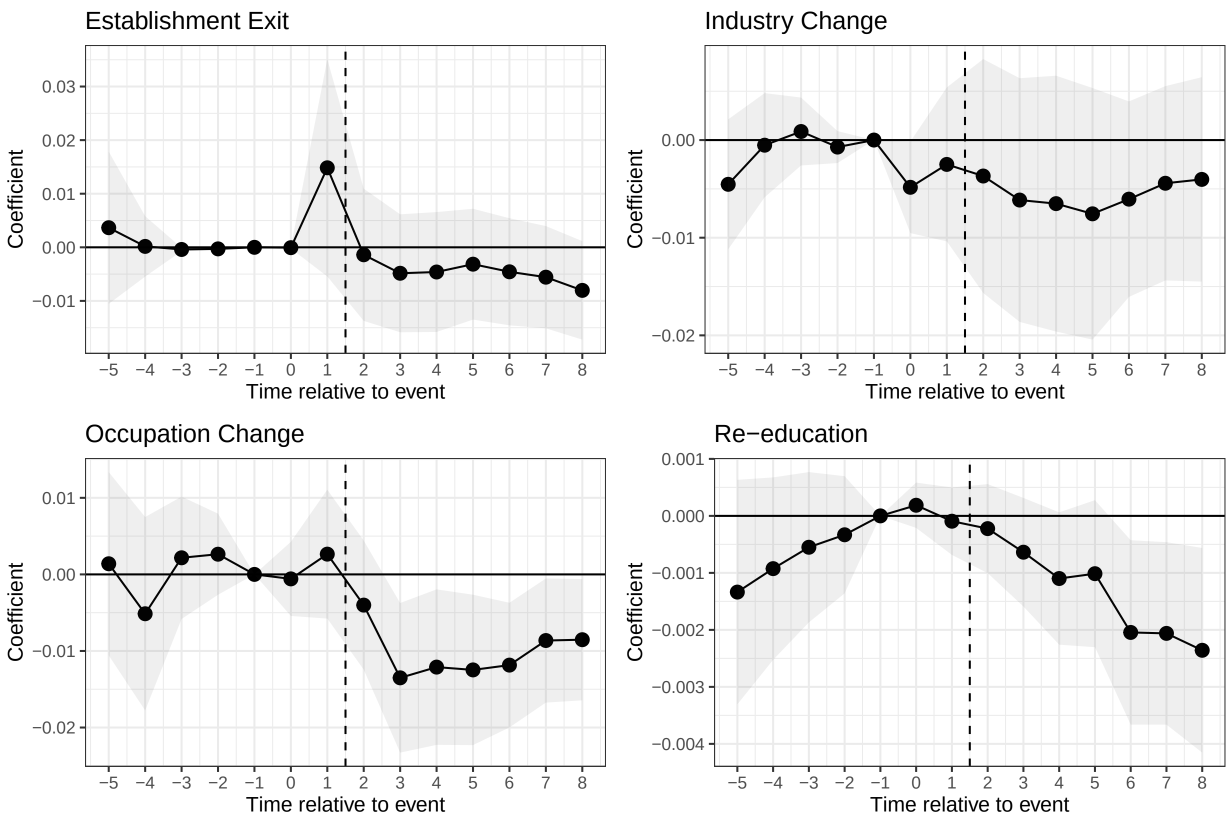

Figure 2-8: Heterogeneous Responses by Education: Adaptive Behaviors.

Notes: Each point is a δ_t coefficient from regression 2.4 where Years of Education is used in place of Trait_i. Years of Education is constructed by mapping degrees to their official length (e.g., a master's degree equals 17 years of education). The model is estimated without any of the factor variables. The outcome used in the estimation is indicated in the panel name. All outcomes are binary and measured relative to their period 0 value. Re-education takes the value of 1 if the degree does not match the period 0 degree. Industry and occupation are measured only for the employed, which restricts the estimation sample to those employed in the post-period.

Table 2.1: Cross-Correlations: Main Variables.

	Extraversion	Conscientiousness	Cognitive Ability	Years of Education	School GPA	Earnings at 35−38
Extraversion	1	0.42	0.35	0.25	0.26	0.12
Conscientiousness	0.42	1	0.23	0.28	0.32	0.1
Cognitive Ability	0.35	0.23	1	0.51	0.64	0.17
Years of Education	0.25	0.28	0.51	1	0.61	0.24
School GPA	0.26	0.32	0.64	0.61	1	0.18
Earnings at 35−38	0.12	0.1	0.17	0.24	0.18	1

Notes: Each number is a pairwise correlation coefficient with a person as the unit of observation. Psychological variables and the school GPA are normalized to have a mean 0 and a standard deviation 1 within cohorts. Earnings are recorded by the tax authorities and measured by averaging total labor and entrepreneurial income earned at age 35–38. Years of Education is constructed by mapping degrees to their official length (e.g., a master's degree equals 17 years of education).

Table 2.2: Cross-Sectional Evidence on Earnings.

Dependent Variable:	log(Earnings)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Extraversion	0.242				0.101	0.096
	(0.003)				(0.004)	(0.004)
Conscientiousness		0.201			0.089	0.023
		(0.003)			(0.004)	(0.003)
Cognitive Ability			0.353		0.297	0.121
			(0.003)		(0.004)	(0.004)
Years of Education				0.210		0.158
				(0.001)		(0.002)
Outcome mean	9.85	9.85	9.85	9.82	9.85	9.86
Fixed-effects						
Birth Year (18)	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	476,195	476,195	476,195	500,123	476,195	474,110
R^2	0.01606	0.01187	0.03129	0.05932	0.03653	0.06123
Within R^2	0.01351	0.00931	0.02878	0.05641	0.03403	0.05876

Notes: Each column reports the OLS regressions results from Equation 2.1 with log earnings as the outcome. The unit of observation is the person. Extraversion, conscientiousness, and cognitive ability are constructed using exploratory factor analysis and normalized to have mean 0 and standard deviation 1 within cohorts. Years of education is constructed by mapping the highest degree at age 35 to its official length (e.g., a high-school degree equals 12 years of education). Earnings are measured by averaging total labor and entrepreneurial income earned at age 35–38.. Heteroskedasticity-robust standard-errors are in parentheses.

Table 2.3: Cross-Sectional Evidence on Adaptive Behaviors.

Dependent Variables:	Total emp.	Occupations	Industry	Establishments
Model:	(1)	(2)	(3)	(4)
Variables				
Extraversion	0.056	0.105	0.070	0.122
	(0.003)	(0.004)	(0.002)	(0.002)
Cognitive Ability	0.313	0.143	0.012	0.061
	(0.004)	(0.004)	(0.002)	(0.002)
Conscientiousness	0.189	-0.016	-0.078	-0.101
	(0.003)	(0.004)	(0.002)	(0.002)
Outcome mean	9.89	2.17	1.56	2.10
Fixed-effects				
Birth Year	Yes	Yes	Yes	Yes
Fit statistics				
Cohorts available	18	4	18	18
Observations	479,820	101,742	479,820	479,820
R^2	0.05448	0.03032	0.01975	0.02430
Within R^2	0.04079	0.02950	0.00581	0.00897

Notes: Each column reports the OLS regressions results from Equation 2.1 with different outcomes. The unit of observation is the person. Total employment is the number years employed at age 28–38. Occupations, Industries, and Establishments represent the total number of different occupation/industry/establishment codes that the individual has worked in at age 28–38. Extraversion, conscientiousness, and cognitive ability are constructed using exploratory factor analysis and normalized to have mean zero and standard deviation 1 within cohorts. Heteroskedasticity-robust standard-errors are in parentheses.

Table 2.4: Baseline Difference-in-Differences Estimates.

Dependent Variables:	Earnings	Employment
Model:	(1)	(2)
Variables		
Post	-0.0129	-0.0258
	(0.0022)	(0.0015)
Post × Treat	-0.0987	-0.0617
	(0.0074)	(0.0043)
Outcome mean	1	0.9700
Fixed-effects		
Person×Event Year (82,405)	Yes	Yes
Age (26)	Yes	Yes
Fit statistics		
Observations	1,349,627	1,349,627
R^2	0.43976	0.29319
Within R^2	0.00811	0.00818

Notes: Each column reports the OLS regression results from Equation 2.3 with different outcomes. The unit of observation is the person-year. Earnings are measured by dividing total labor income with period 0 earnings. Employment is binary and takes the value 1 if the individual is employed during the last week of the year. The post-period indicator includes 7 years after the event and 5 years before the event. The event year is omitted from the estimation sample. One-way (Establishment) standard-errors are in parentheses.

Table 2.5: Triple-Difference Estimates: Earnings.

Dependent Variable:	Earnings			
Model:	(1)	(2)	(3)	(4)
Variables				
Post × Treat × Extraversion	0.021		0.021	0.018
	(0.004)		(0.004)	(0.004)
Post × Treat × Conscientiousness	0.0009		-0.0003	0.001
	(0.004)		(0.004)	(0.004)
Post × Treat × Cognitive Ability	0.009		0.005	0.007
	(0.005)		(0.004)	(0.005)
Post × Treat × Age	-0.003	-0.003	-0.003	-0.002
	(0.001)	(0.001)	(0.001)	(0.001)
Post × Treat × Years of Education		0.008	0.004	0.003
		(0.004)	(0.004)	(0.002)
Outcome mean	1	1	1	0.990
Fixed-effects				
Event Year×Person	Yes	Yes	Yes	Yes
Post×Treat×Occupation (172)				Yes
Post×Treat×Industry (482)				Yes
Fit statistics				
Event Year×Person	82,405	82,405	82,405	57,129
Observations	1,349,627	1,349,627	1,349,627	945,820
R^2	0.44399	0.44532	0.44700	0.45732
Within R^2	0.09100	0.09317	0.09591	0.03468

Notes: Each column reports the OLS regression results from Equation 2.5 with earnings as the outcome. The unit of observation is the person-year. Earnings are measured by dividing total labor income with period 0 earnings. Extraversion, conscientiousness, and cognitive ability are constructed using exploratory factor analysis and normalized to have mean zero and standard deviation 1 within cohorts. The post-period indicator includes 7 years after the event and 5 years before the event. The event year is omitted from the estimation sample. One-way (Establishment) standard-errors are in parentheses.

Table 2.6: Triple-Difference Estimates: Employment.

Dependent Variable:	Employment			
Model:	(1)	(2)	(3)	(4)
Variables				
Post × Treat × Extraversion	0.013		0.013	0.010
	(0.002)		(0.002)	(0.003)
Post × Treat × Conscientiousness	0.0005		-0.0008	0.0005
	(0.002)		(0.002)	(0.003)
Post × Treat × Cognitive Ability	0.013		0.008	0.007
	(0.003)		(0.003)	(0.003)
Post × Treat × Age	-0.001	-0.001	-0.001	-0.0006
	(0.0007)	(0.0007)	(0.0007)	(0.0007)
Post × Treat × Years of Education		0.008	0.004	0.004
		(0.002)	(0.002)	(0.002)
Outcome mean	0.970	0.970	0.970	0.960
Fixed-effects				
Event Year×Person	Yes	Yes	Yes	Yes
Post×Treat×Occupation (172)				Yes
Post×Treat×Industry (482)				Yes
Fit statistics				
Event Year×Person	82,405	82,405	82,405	57,129
Observations	1,349,627	1,349,627	1,349,627	945,820
R^2	0.29543	0.29496	0.29586	0.30855
Within R^2	0.05453	0.05391	0.05511	0.00670

Notes: Each column reports the OLS regression results from Equation 2.5 with employment as the outcome. The unit of observation is the person-year. Employment is binary and takes the value 1 if the individual is employed during the last week of the year. Extraversion, conscientiousness, and cognitive ability are constructed using exploratory factor analysis and normalized to have mean zero and standard deviation 1 within cohorts. The post-period indicator includes 7 years after the event and 5 years before the event. The event year is omitted from the estimation sample. One-way (Establishment) standard-errors are in parentheses.

Table 2.7: Triple-Difference Estimates: Adaptive Behaviors.

Dependent Variables:	Occupation	Industry	Education	Establishment Exit
Model:	(1)	(2)	(3)	(4)
Variables				
Post × Treat × Extraversion	-0.017	-0.012	-0.004	-0.016
	(0.006)	(0.005)	(0.002)	(0.005)
Post × Treat × Conscientiousness	0.009	-0.0001	0.001	0.006
	(0.006)	(0.005)	(0.002)	(0.005)
Post × Treat × Cognitive Ability	-0.011	-0.008	-0.003	-0.002
	(0.008)	(0.009)	(0.002)	(0.008)
Post × Treat × Age	-0.0005	0.004	0.0009	-0.0005
	(0.002)	(0.002)	(0.0005)	(0.001)
Outcome mean	0.290	0.130	0.030	0.220
Fixed-effects				
Event Year×Person	Yes	Yes	Yes	Yes
Fit statistics				
Event Year×Person	57,129	82,405	82,405	82,405
Observations	761,234	1,289,243	1,349,627	1,349,627
R^2	0.50543	0.44616	0.41160	0.51031
Within R^2	0.17526	0.11556	0.01644	0.26900

Notes: Each column reports the OLS regression results from Equation 2.5 with different outcomes. The unit of observation is the person times year. All outcomes are binary and measured relative to their period 0 value. Education takes value 1 if the degree does not match the period 0 degree. Industry and occupation are measured only for the employed, which restricts the estimation sample to those who are employed in the post-period. Extraversion, conscientiousness, and cognitive ability are constructed using exploratory factor analysis and normalized to have mean zero and standard deviation 1 within cohorts. The post-period indicator includes 7 years after the event and 5 years before the event. The event year is omitted from the estimation sample. One-way (Establishment) standard-errors are in parentheses.

Table 2.8: Triple-Difference Estimates: Adaptive Behaviors With Additional Controls.

Dependent Variables:	Occupation	Industry	Education	Establishment Exit
Model:	(1)	(2)	(3)	(4)
Variables				
Post × Treat × Extraversion	-0.013	-0.003	-0.0002	-0.005
	(0.006)	(0.006)	(0.002)	(0.005)
Post × Treat × Conscientiousness	0.011	-0.006	0.0002	0.0008
	(0.005)	(0.005)	(0.003)	(0.005)
Post × Treat × Cognitive Ability	0.004	-0.011	-0.003	-0.002
	(0.007)	(0.006)	(0.003)	(0.006)
Post × Treat × Age	0.0005	0.003	0.001	0.002
	(0.002)	(0.002)	(0.0006)	(0.002)
Post × Treat × Years of Education	-0.005	0.0005	0.002	-0.006
	(0.003)	(0.004)	(0.001)	(0.003)
Outcome mean	0.290	0.130	0.030	0.210
Fixed-effects				
Post×Treat×Occupation (172)	Yes	Yes	Yes	Yes
Post×Treat×Industry (482)	Yes	Yes	Yes	Yes
Event Year×Person (57,129)	Yes	Yes	Yes	Yes
Fit statistics				
Observations	761,234	904,617	945,820	945,820
R^2	0.52471	0.48735	0.45917	0.53177
Within R^2	0.01026	0.00574	0.07437	0.01516

Notes: Each column reports the OLS regression results from Equation 2.5 with different outcomes. The unit of observation is the person times year. All outcomes are binary and measured relative to their period 0 value. Education takes value 1 if the degree does not match the period 0 degree. Industry and occupation are measured only for the employed, which restricts the estimation sample to those who are employed in the post-period. Extraversion, conscientiousness, and cognitive ability are constructed using exploratory factor analysis and normalized to have mean zero and standard deviation 1 within cohorts. The post-period indicator includes 7 years after the event and 5 years before the event. The event year is omitted from the estimation sample. One-way (Establishment) standard-errors are in parentheses.

Chapter 3

School vs. Action Oriented Traits in the Labor Market

With Ramin Izadi

3.1 Introduction

Extensive evidence shows that noncognitive skills[1] improve labor market success (Almlund et al., 2011; Deming, 2017), but the channel is incompletely understood. Some studies show that noncognitive skills affect labor-market performance indirectly through higher educational aptitude (Cunha and Heckman, 2007), while other studies emphasize that noncognitive skills affect labor productivity directly at work (Deming, 2017).

How do different dimensions of personality predict school vs. labor-market performance? How has the labor-market value of these traits changed over time? We answer these questions using globally exceptional data that includes multidimensional personality and cognitive test scores, education, and labor-market records for 79% of Finnish men born 1962–1979 (n = 489,252). The personality and cognitive test data were collected by the Finnish Defence Forces during mandatory military service.

This paper shows that some dimensions of noncognitive skills are productive at school and also valued in the labor market, while other dimensions are counterproductive at school yet still valued in the labor market. We further document that the labor-market returns to action-oriented personality traits (traits that predict low school performance) have rapidly increased over the past two decades. Conversely, the economic returns to school-oriented traits have declined sharply.

Consider the school versus the labor-market. Noncognitive skills related to conscientiousness have been shown to predict school achievement (Almlund et al., 2011). At the same time, the

[1] Noncognitive skills are typically defined as all skills not predicted by cognitive test scores. In some contexts, noncognitive skills specifically refer to socioemotional skills. We adopt the standard definition that noncognitive skills refer to all potentially economically valuable traits that the cognitive tests do not measure. In this view, we define some personality traits as noncognitive skills.

common stereotypes of socially awkward 'nerds' and outgoing 'jocks' suggest an inverse relationship between school achievement and particular dimensions of noncognitive skills.[2] High achievers in school may lack at least in some dimensions of economically valuable personality traits, and conversely, low-achieving students may have some redeeming qualities that compensate in the labor market for their lack of academic success. In short, this idea suggests a negative association between academic performance and outward-oriented social skills.

We present four new descriptive facts. First, we document that one subset of personality traits positively predicts school achievement, but another critical subset of personality traits negatively predicts school achievement. These subsets follow the common stereotypes: men, who score highly in activity-energy, sociability, and masculinity, tend to perform worse in standardized tests. We label this component as action-oriented traits. The label also reflects the source of measurement: The Finnish Defence Forces values these traits positively. In contrast, dutifulness, deliberation, achievement striving, self-confidence, and leadership predict good school performance; we label this component as school-oriented traits.

Second, the traits that predict low school achievement still predict labor-market success. One standard deviation increase in action-oriented traits predicts a 5-log point increase in earnings at age 35. The school-oriented traits also strongly predict labor-market success. But the school-oriented traits are not independently valued in the labor market: their predictive power on labor-market performance becomes near zero after controlling for school achievement.

Third, we find that the labor-market returns to action-oriented traits that predict low school performance have rapidly increased over the past 17 years, from 0 to 8 log points per standard deviation. Men with high activity-energy, sociability, and masculinity measures (but with low mathematics skills) experienced the highest earnings gains between 1997 and 2017. The returns to school-oriented traits have declined from 13 to 7.5 log points.

Fourth, specialization of skills has increased over the past two decades: men are more likely to have either high mathematics skills or action-oriented traits and are less likely to have both.

To understand the empirical results, we outline a model of multidimensional skill specialization. Intuitively, the model highlights two paths to labor-market success: one through school-oriented traits and formal skills, and one through action-oriented traits and informal skills. In the model, the labor market rewards individuals for their formal skills gained through education and for their informal skills, e.g., initiative, social skills, and charisma. Personality is a fixed endowment for an individual, but skills are endogenous and require a time investment. At the investment stage, individuals can allocate their time between study and activities that improve their informal skills, such as social life. We model personality by two separate dimensions: traits that increase productivity in informal-skill formation (action-oriented) and traits that make studying more efficient (school-oriented). Heterogeneity in the initial endowment of traits generates a comparative advantage in formal or informal skill accumulation. In equilibrium, this comparative advantage drives individuals to specialize relatively more in the type of human capital where they have pre-existing

[2]Stereotype accuracy is one of the most replicable findings in social psychology (Jussim et al., 2016). Both stereotypes we mention can have negative connotations when referring to a person.

tendencies.

We interpret the findings using the model. First, we demonstrate that both the action-oriented and school-oriented personality traits have a positive return in an earnings regression. But controlling for standardized test scores, the return for action-oriented traits increases, and the return for school-oriented traits becomes small. This pattern arises from the intransitivity of correlations between action-oriented traits, test scores, and adult earnings. Our model rationalizes this intransitivity: higher endowment in action-oriented traits increases investment in informal skills at the expense of school success. Since test-score performance is endogenous, its inclusion inflates the returns to action-oriented traits and deflates the returns to school-oriented traits. Our model allows traits to directly affect earnings beyond their instrumental effects through informal and formal skills (e.g., education). Looking through our model, the low returns to school-oriented traits when test scores are included suggest that their effect would mostly be mediated by educational achievement. We cannot similarly disentangle the direct effects of action-oriented traits from the returns to informal skills because we cannot directly measure informal skills. Therefore, in our empirical work, a single variable captures both the direct effect of action-oriented traits and the indirect effects through informal skills.

Next, we explore the channels through which the returns are realized in the labor market by estimating a model where personality traits and test scores explain different response variables. The traits that predict high school achievement (school-oriented traits) appear to affect labor-market performance mainly through occupational sorting, and the traits that predict low school achievement (action-oriented traits) primarily through within-occupation effects and work experience. On average, action-oriented individuals are not more likely to select into high-paying occupations. Instead, action-oriented individuals acquire less education and start their careers earlier but with fewer unemployment spells. Specifically, they are less likely to select into high-paying professional occupations, typically not available without higher education. But even with lower education, they are more likely to end up in a managerial position. In contrast, individuals with high school-oriented traits are more likely to select into high-paying professional occupations. They acquire higher education, start their careers later, and spend less time in unemployment. Occupational and educational sorting explains a large part of earnings variations for both types of individuals. But when including fixed effects for education and occupation, action-oriented traits become a significantly larger predictor of earnings than mathematics. In total, we interpret this as evidence that action-oriented traits improve earnings, mainly through experience, job performance, and/or career progress. In contrast, personality that predicts higher educational attainment helps individuals start their careers in higher-paying jobs but plays a smaller role afterward.

Finally, we document two novel time trends. First, the return to action-oriented traits has increased markedly during our 17-year measurement period. The finding is consistent with earlier studies on the returns to social skills (Deming, 2017) and non-cognitive skills (Edin et al., 2021). But it also provides a new angle: We observe a similar rise to other action-oriented traits---activity and masculinity. At the same time, the returns to traits that predict school-performance have declined. Second, specialization into two distinct types has become more common: more students have either low mathematics and high action-oriented traits *or* high mathematics and low action-oriented traits. In our model, these trends are consistent with a supply-side response, where the increasing returns to informal skills reinforce the skill specialization.

This paper contributes to several distinct lines of research.

Multidimensional skills. Re-emerging literature highlights the importance of the multidimensional nature of skills. The idea that skills are multidimensional is not new. For example, the classic Roy (1951) model formalizes the idea that workers may differ in their types of skills (hunting vs. gathering), and that this affects the optimizing choices of workers selecting between bundles of tasks. Gardner (1983) differentiates skills into specific 'modalities,' imperfectly described by a unidimensional skill.

An emerging line of economic research, both empirical and theoretical, focuses on the multidimensional match between skills and tasks (Guvenen et al., 2020; Lise and Postel-Vinay, 2020; Fredriksson et al., 2018; Lindenlaub, 2017; Groes et al., 2015; Gathmann and Schönberg, 2010). These studies emphasize the potential for skill mismatch: a situation where worker's bundle of skills is not well-matched with the distribution of skill-requirements for the set of tasks.

This paper provides novel descriptive facts to advance this literature. It illustrates skill specialization in the supply side of multidimensional skills. We document that this specialization has concrete implications on occupational sorting and earnings.

Noncognitive skills. A large literature analyzes the role of noncognitive skills in the labor market. The evidence unambiguously demonstrates that a wide array of noncognitive skills—personality traits, interpersonal skills, etc.—are important drivers of labor-market success (Heckman et al. 2006; Lindqvist and Vestman 2011; Weinberger 2014; Deming 2017; Jokela et al. 2017).[3] The research on noncognitive skills emphasizes two channels on how noncognitive skills may affect labor-market performance: the direct channel, for example, social skills facilitating teamwork in production (Deming 2017), and the indirect channel, for example, noncognitive skills fostering cognitive skills and human-capital production (Cunha and Heckman 2007; Borghans et al. 2016).

We complement this literature by showing that the noncognitive skills associated with the indirect channel are notably different from those associated with the direct channel. That is, the traits that predict good school performance are different from those that predict good labor-market performance (e.g. conscientiousness vs. extraversion). In particular, we show that social skills—while important in the labor market, e.g., for teamwork (Deming, 2017)—are negatively correlated with academic test-scores.

In this line of work, the most closely related research are Levine and Rubinstein (2017) who show that the inverse 'combination of 'smart' and 'illicit' tendencies as youths' predict entry into and success in entrepreneurship, Papageorge et al. (2019) who argue that some childhood misbehavior represents socio-emotional skills that are valued in the labor market, Lleras-Muney et al. (2020)

[3] Almlund et al. (2011) provide an excellent survey of the evidence on the predictive power of personality in the labor market.

who analyze the dual decision of investment in education and social capital, and Bursztyn et al. (2019) who highlight the trade-off between social-image concerns and school effort.

Education. Our findings provide an explanation to the 'reading penalty paradox' documented by Altonji et al. (2016) and Sanders (2015). Frequently-used US data sets that include information on test scores and earnings exhibit a negative association between reading scores and earnings once the researchers control for mathematics test scores. This pattern also arises in our data. But it goes away once we control for personality traits. The predictive returns to verbal skills are close to zero with personality controls. These pieces of evidence suggest that the observed 'reading penalty' emerges from omitted economically valuable noncognitive skills that are negatively correlated with verbal skills.

Trends in returns to skills. Long-standing literature estimates the returns to skills over time (Katz and Murphy 1992; Goldin and Katz 2008; Acemoglu and Autor 2011; Deming 2017; Edin et al. 2021). We show that the returns to those skills that predict low school performance, sociability, activity, and masculinity, have rapidly increased over the past 17 years in Finland. At the same time, the returns to cognitive skills has been remarkably stable. In the supply side, we show that the skill specialization, to school-oriented traits and formal skills and to action-oriented traits and informal skills, has increased over the past 17 years.

3.2 Data

This project combines several data sources using unique person identifiers.[4]

Personality and Cognitive Skills Data for personality and cognitive skills are obtained from the *Finnish Defence Forces* (FDF), which has tested all military conscripts since 1955. The available data cover 79% of Finnish men born between 1962 and 1979 (n = 489,252).

The data provide detailed test scores for personality (8 dimensions) and cognitive skills (3 dimensions). The measured personality traits are: sociability, activity-energy, masculinity, dutifulness, deliberation, achievement motivation, leadership motivation, and self-confidence. The measured cognitive skills are visuospatial, arithmetic, and verbal reasoning. The visuospatial test is similar to Raven's Progressive Matrices (Raven and Court, 1938).

The personality dimensions are based on the Minnesota Multiphasic Personality Inventory (MMPI) which predates the Big Five model by several decades. That is why it includes a somewhat different set of items compared to the Big Five inventory. However, two of the Big Five traits are represented by their facets (subtraits). Dutifulness and deliberation are subtraits associated with conscientiousness, whereas the subtraits sociability, activity-energy and self-confidence are associated with extraversion. Masculinity is not measured in all standard personality inventories but turns out to be an important predictor in our analysis.

Military conscription in Finland is universal and grants relatively few exceptions. Finnish men

[4]The data are described in more detail in Appendix C.1.

are drafted in the year they turn 18 and most start their service at age 19 or 20. Military service lasts for 6–12 months, and most conscripts do not continue service at the military. FDF uses psychological tests to assess conscripts' suitability for non-commissioned officer training.

Both personality and cognitive ability tests are typically taken in the second week of military service in a 2-h paper-and-pencil format in standardized group-administered conditions. The personality test contains 218 statements with a response scale of yes/no. The cognitive test contains 120 multiple-choice questions. The test questionnaires have been unchanged for the timeline of the study, and the scores are designed to be comparable across cohorts. Appendix C.1 includes a more detailed description of the FDF data.

Education Data on education come from three sources.

The Register of Completed Education and Degrees contains exact information on the educational degrees the individual has obtained, including both the level and field, and the date at which the degree was granted. All degrees completed in Finland are generally recorded in these data.

The Secondary Education Application Register contains information on the 9th-grade transcript, including the GPA. The data are produced as a side product of the centralized application system for secondary education maintained by The Finnish National Board of Education (FNBE). While the 9th-grade records are only partly from national standardized tests, the middle schools in Finland are all public and have low quality variance. Attendance of the 9th grade is near-universal. These data are only available for cohorts born 1975-1979.

Finnish Matriculation Examination Board Register (FMEB) contains test-score data by academic subject in the standardized national-level high-school exit examination, The Matriculation Examination (ME). Independent reviewers grade the test in a double-blind manner, and within the timeline of this study, the test scores directly correspond to ranks within a subtest and cohort. The students choose a minimum of four 6-hour tests in their first language, foreign language, mathematics, and in the subjects of humanities and natural sciences. The first-language test is mandatory. Language and mathematics tests have basic and advanced-level versions. When needed for the analysis, we map the mathematics test scores into a single dimension by weighting the advanced and basic test scores using their predictive power on the military arithmetic test.[5] As an institutional background, secondary schooling in Finland has two tracks: academic and vocational. Participation to the academic track increased from 35% to 47% between birth cohorts 1962–1979. ME is the academic track's exit exam. A similar standardized test does not exist for the vocational track. ME test scores are partly used in university admissions (most university admissions, however, were based on a separate admissions exam), but they do not play a meaningful independent role in the labor market. The Finnish school system contains relatively few extracurricular activities (such as sports teams) that could be used to measure non-school human capital, and the participation or performance in these activities are generally not recorded.

[5] This procedure is described in more detail in AppendixC.1.

Labor Market and Demographics The project uses detailed longitudinal register data on the full Finnish population compiled by *Statistics Finland* from multiple sources.

The register data provide information on demographics, labor market status, earnings, occupation, industry, firm and establishment identifiers, county of residence and birth, and the identity of parents and siblings, for all Finnish residents.

Income data are obtained from the *Finnish Tax Authority*. We measure 'prime-age' earnings as the average annual labor-market earnings during ages 35–38. We deflate all values to 2010 Euros using the Statistics Finland CPI.

Main Estimation Sample Our main estimation sample consists of the intersection of individuals with valid (1) military test scores, (2) high school exit exam records, and (3) positive prime age earnings (over 99%). The sample size is approximately 158,000 containing about 80% of male high school graduates born in 1962–1979. Potential selection issues are discussed in Appendix C.1.

For our main analysis, we use logarithmic earnings. Figure 3-1 shows their distribution in our main sample. The long left tail typical for log earnings distributions can raise concerns about outliers driving our OLS results. These concerns are addressed in section 3.4 and Appendix C.3.

3.3 Model

In this section, we develop a simple model of multidimensional skill specialization that provides a structure for the relationships between personality, education, and labor-market performance.[6] We focus on the distinction between the production of human capital and the productive activities in the labor market. The personality traits are viewed as a fixed type and skills are viewed as endogenous. For concreteness, the context can be thought as students in high school and the labor-market.

At the center of the model, there are two production functions for two types of human capital, formal skills ('education', H) and informal skills ('social capital', S):[7]

$$H(h; N, J) = a(N, J) \times h \tag{3.1}$$

$$S(s; N, J) = b(N, J) \times s \tag{3.2}$$

Formal skills H are produced by time investment h ('studying') and informal skills S are produced by time investment s ('socializing'). The productivities of human capital production, a and b, depend on the endowment of personality traits (N for 'school-oriented' and J for 'action-oriented').

In making a decision, the students face a time-allocation constraint:

$$h + s = T. \tag{3.3}$$

[6]The model is related to Lleras-Muney et al. (2020).

[7]For simplicity, we consider two 'technologies of skill formation' (Cunha and Heckman, 2007). The case for n technologies is analogous.

Time spent on studying is away from time spent on socializing. We normalize the time endowment as $T = 1$.

The objective function is:

$$U(s; N, J) = H(1 - s; N, J) + S(s; N, J) + V(s; N, J) \tag{3.4}$$

The students value both types of human capital, H and S, but also derive direct utility (or disutility) from studying and socializing, V, that depends on their endowment of traits.[8] We further assume that the direct utility function depends only on the relative allocation between h and s.[9]

Students choose how much time to spend on socializing (s) to maximize the objective function under the constraint. With no strategic behavior or dynamics, the optimal time-allocation decision between studying and socializing is a static individual optimization problem.[10] The first-order condition for socializing is:

$$\frac{\partial V(s; N, J)}{\partial s} = a(N, J) - b(N, J) \tag{3.5}$$

$$s^*(N, J) = g_s(a(N, J) - b(N, J); N, J) \tag{3.6}$$

where g_s is the inverse function of $\partial V(s, N, J)/\partial s$ with respect to s. For an interior solution to exist, Expression 3.5 must be positive. Intuitively, if at the optimal s^*, socializing is not only more fun ($\partial V(s; N, J)/\partial s > 0$) but also more productive ($b(N, J) > a(N, J)$), there would be no reason to study at all.

The theoretical analysis focuses on the decision to socialize, s; the analysis for the inverse decision of studying, $h = 1 - s$, is symmetric. We provide proofs in Appendix C.2.

This setup provides some flexibility by admitting at least three distinct interpretations. In the classic view, students gain utility from formal skills H and informal skills S because there is a return to different human-capital types in the labor market. Students also face a direct cost or benefit from studying and socializing V that depends on their endowments. In this view, socializing is an investment: students socialize not just because studying may be laborious but also because socializing builds people-skills and networks rewarded in the labor market.

From a more modern perspective (see, for example, Lavecchia et al. 2016), students might not be sufficiently forward-looking to consider their future earnings. However, the terms H and S can be interpreted as social norms that guide their choices, for example, through parental pressure.

[8] If U reflects log utility, the objective function arises under the canonical CES preferences. At this point, to keep the notation clear, we abstract from potential return multipliers for H and S in the objective function.

[9] Because time is spent between studying and socializing, V captures all direct costs/benefits of studying/socializing. For example, action-oriented students could dislike studying but school-oriented students might enjoy it.

[10] The model has an implied timing that corresponds to a typical path from adolescence to adulthood. Students enter a schooling period with an initial personality endowment (N, J). They then decide how much time to spend on socializing s. Their H and S are realized at the end of the schooling period. After the schooling period, they enter the labor market and receive earnings Y.

In this interpretation, the cost function V is the direct utility of socializing over and above the social-norm component S.[11]

Finally, we could abstract entirely from the source of utility derived from either type of human capital *stock*. Students simply enjoy the activity of spending time s with their friends. Performing well in tests requires time to study (h), which may be uninteresting and incurs a cost $-V$. From this perspective, V reflects the direct utility of time spent socializing, which may be different for students with different personality endowments.

Skill Specialization The fundamental trade-off between time investments in this model leads to skill specialization, where students with a comparative advantage in the action-oriented endowment invest more time on socializing relative to students with a comparative advantage in the school-oriented endowment. Taking a derivative of the first order condition in Equation 3.5 and solving for $\partial s^*(N,J)/\partial J$ gives:

$$\frac{\partial s^*(N,J)}{\partial J} = -\left[\underbrace{\frac{\partial^2 V(.)}{\partial s^2}}_{<0}\right]^{-1}\left[\underbrace{b^J(N,J) + \frac{\partial^2 V(.)}{\partial s \partial J}}_{\text{marginal benefit of } \Delta J} - \underbrace{a^J(N,J)}_{\text{marginal cost of } \Delta J}\right] > 0. \tag{3.7}$$

The first term is the gradient in the marginal direct utility of socializing (or marginal cost of studying). We assume the standard decreasing marginal utility. Hence, for socializing: $\partial^2 V(.)\big/\partial s^2 < 0$. We also assume that the productivity of informal-skill accumulation $b(N,J)$ is increasing in the action-oriented trait-endowment J. This assumption is based on the idea that learning social skills, creating networks, and improving their social hierarchy position is easier for students who already have sociable and proactive personalities. Likewise, we assume that action-oriented students enjoy a larger marginal utility of socialization s. Formally, $\partial^2 V(.)\big/\partial s \partial J > 0$. This reflects the idea that the opportunity cost for studying is higher for action-oriented students who could be having fun with their friends instead. Finally, we assume that productivity of studying, $a(N,J)$, does not depend on the action-oriented trait, so that $a^J(N,J) = a^J(N) = 0$. With these key assumptions, the right-hand side of Equation 3.7 is positive, and an increase in the action-oriented endowment leads to an increase in the time spent socializing.

Immediately following from these assumptions we also have:

$$\frac{\partial H(1-s;N,J)}{\partial J} = \frac{\partial a(N)(1-s^*(N,J))}{\partial J} = -a(N)\frac{\partial s^*(N,J)}{\partial J} < 0. \tag{3.8}$$

In other words, conditional on the school-oriented trait N, more action-oriented students have worse test scores. It implies that comparative advantage determines the time-allocation decision.

Figure 3-2 simulates the model with a quadratic cost function and linear productivity functions. Each line represents an isoquant where the optimal time allocation decision s^* does not change.

[11] The cost of studying and the direct utility of socializing mirror each other, because students allocate time T between the two.

Along each line, as long as the *comparative* proportion of endowments does not change, the *absolute* levels can vary substantially, still resulting in the same optimum allocation. For example, at the bottom right-hand corner, investment in s is highest; these are students who are high in J but low in N. At the upper left-hand corner are students high in N but low in J; their investment in s is lowest.

Returns to Personality The labor market rewards both types of human capital, H and S, and also directly the endowments, N and J. Earnings are determined by:

$$Y = r_H H(1-s;N,J) + r_S S(s;N,J) + r_N N + r_J J \tag{3.9}$$

$$= r_H a(N)(1-s^*(N,J)) + r_S b(N,J) s^*(N,J) + r_N N + r_J J \tag{3.10}$$

where r_H and r_S are the returns to the respective dimensions of human capital and r_N and r_J are the direct returns to the respective traits.

The marginal returns to the action-oriented trait are:

$$\frac{\partial Y}{\partial J} = \underbrace{r_S b^J(J,N)s^*(N,J) + r_J}_{\text{direct effect of } \Delta J} + \underbrace{(r_S b(J,N) - r_H a(N))}_{\text{net earnings change for } \Delta s} \underbrace{\frac{\partial s^*(N,J)}{\partial J}}_{\Delta s} \lesseqgtr 0. \tag{3.11}$$

The first term is the direct effect: the effect of the increase in the productivity of informal-skill production and the direct return from the increase in the action-oriented trait. By assumption, productivity $b(J)$ is increasing in the action-oriented trait, so this term is positive. The second term is the indirect effect: the change in earnings due to changes in the optimal time allocation s^*. As shown earlier, an increase in the action-oriented trait leads to an increase in the time investment (Δs). This reallocation results in a shift from formal skills to informal skills. However, at the optimum, as shown in Expression 3.5, the productivity of informal skills must be lower than the productivity of formal skills. Taken together, the indirect effect is negative.

Intuitively, students with a lower initial comparative advantage in socializing (low action-oriented trait in relative terms) take a larger hit from investing more in s, because their comparative advantage is in formal skills (or educational capital), from which they are substituting away by increasing s. At the same time, $\partial s^*(N,J)/\partial J$ is smaller for students with a comparative advantage in studying.

In total, the sign of Expression 3.11 is ambiguous. If the gains from the direct effect are larger than the losses from the indirect effect, we should expect a positive return to the action-oriented trait conditional on the school-oriented trait.

3.4 Results

3.4.1 Personality and Academic Performance

This section establishes the empirical relationship between academic achievement and our personality measures and uses the results to create a simple 'trait taxonomy' for subsequent analysis. We estimate:

$$H_{it} = P_i'\beta + \delta_t + \epsilon_{it} \tag{3.12}$$

where H_{it} is a test score measuring academic performance and P_i is a vector of personality traits for person i in birth cohort t.

Table 3.1 shows the OLS regression results, where the eight personality traits are used as linear predictors for the high-school test scores. Each test score and personality trait is normalized within the analyzed sample. The first four columns represent different high-school subjects as outcome variables. A clear pattern emerges from the partial correlations: sociability, activity-energy, and masculinity consistently negatively predict academic test scores, while deliberation, dutifulness, achievement-striving, and leadership motivation positively predict higher test scores.[12] The relative impact varies somewhat by subject, but confidence and achievement-striving are the strongest positive predictors, whereas sociability is the strongest negative predictor, with the exception of the verbal test, where all three negative traits have roughly equal importance. Overall, personality traits explain 9–13% of the variation in high-school test scores.

The pattern holds for a wide set of educational outcomes. Columns 5-7 show that it holds for 9th grade GPA, selection into high school, and years of schooling. At every level, individuals with high sociability, activity-energy, and masculinity place lower in the intensive margin (grades/test scores) as well as in the extensive margin (selection into education).

Next, we use the dichotomy of positive and negative traits to reduce individual personality into just two distinct dimensions: an index that predicts positive test-score performance and an index that predicts negative test-score performance. For each individual, each index takes a value that is the weighted average of the corresponding traits. Column 4 of Table 3.1 shows the weights from the anchoring regression, where the overall test score average is used as the response variable. While the choice of the response variable is somewhat arbitrary, our analysis is robust to using any specific test score as the response variable.[13] For the rest of the paper, based on their intuitive appeal, we call the index of positive traits 'school-oriented' and the inverse of the index of negative traits 'action-oriented'. Note that 'positive' and 'negative' are used only to describe their association with educational outcomes. We do not imply that action-oriented traits in this context are negative in

[12]In a misspecified model, this pattern could arise incidentally in the presence of multicollinearity between personality traits. Table C-1 in Appendix C.3 shows a full cross-correlation table that demonstrates the same pattern with pairwise correlations.

[13]An alternative weighting scheme is to use component scores from principal component analysis (separately applied to negative/positive traits). The indices obtained in this manner have a 0.99 correlation with indices obtained from our preferred anchoring method.

any other sense.

Our model provides a framework to understand these results. In the model, some personality traits make social activities more attractive at the expense of studying, and vice versa. This causes students highly endowed on those traits to allocate their time differently, resulting in the inverse bundling of those traits with school performance. The strength of the inverse relationship between school test scores and personality traits depends on the joint distribution of personality traits (N and J in our model) and the impact of personality on the utility of socializing (V) and the productivity of socializing and studying (functions a and b).

In subsequent analysis, we apply a similar anchoring procedure to compare the test scores of the two different tracks of mathematics, basic and advanced. A single mathematics test score for each individual is obtained by regressing the military mathematics test results (available for everyone in the sample) on the test scores from the two high-school mathematics tracks and using the coefficients to weight the test scores. Appendix C.1 outlines the details of this procedure.

3.4.1.1 Why not Use Factor Analysis and the Big Five Model?

The usual approach to dimension reduction with multidimensional data follows the principles of Exploratory Factor Analysis (EFA). The idea is to group closely correlated variables into a single variable---the latent factor. Jokela et al. (2017) and Izadi and Tuhkuri (2021a) conduct EFA using the psychological traits in the Finnish Defence Force data. Izadi and Tuhkuri (2021a) show that reducing the eight personality traits into two factors results in a very different grouping in comparison to the indices used in this paper. Specifically, in EFA, dutifulness and deliberation mostly load onto one factor, and sociability, activity-energy, leadership motivation, and self-confidence load onto the other.[14] This division is broadly consistent with the Big Five model of personality. In the literature, each of the Big Five domains can be further divided into 'facets', or subgroups of traits. Deliberation and dutifulness are facets associated with conscientiousness whereas, sociability, activity-Energy and self-Confidence are facets associated with extraversion.

Importantly, the patterns we find in this paper do not emerge by replacing the school-oriented and action-oriented index by the two factors found with EFA. The reason is that, for example, the factor associated with extraversion assigns positive loadings to both, traits that predict academic success (self-confidence, leadership motivation, achievement motivations), and traits that predict bad academic performance (sociability and activity-energy). In other words, the top level Big Five categorization is too coarse and does not capture the nuanced effects of individual facets/traits adequately for our purpose. The inadequacy of the Big Five domains in predicting behaviors is not itself a new finding (Paunonen and Ashton, 2001; Vainik et al., 2019), but it warrants a method of dimension reduction that is specifically tailored for the behavior that we study: educational attainment. With this in mind, we purposefully group the traits based on their relation to educational attainment, even when they belong to different Big Five domains. The indices that arise from this exercise should not be viewed as factors since they are not constructed based on the

[14] Achievement motivation loads onto both but masculinity does not load strongly onto either.

cross-correlation of individual items with each other, as in factor analysis, but rather on their correlation with a common outcome: educational attainment. For example, masculinity (which is a non-standard item in personality questionnaires), is not strongly correlated with any other trait, but is included in our action-oriented index as a strong predictor of bad school performance. As such, our composite indices are most closely related to the economic concept of 'types', which relate parameters directly to behaviors. This is also reflected in our model. Our approach is in the spirit of recent discussions in personality psychology that emphasize the importance of the narrower facets over the broader domains to explain the causal mechanisms of personality on outcomes (see Mõttus 2016).

3.4.2 Personality and Labor-Market Performance

We estimate an earnings regression where the logarithmic prime-age earnings (Y) are regressed on the intensity of action-oriented (J) and school-oriented (N) traits, and in some specifications also on IQ and high-school test scores (H):

$$Y_{it} = \beta_1 N_i + \beta_2 J_i + H_i' \beta + \delta_t + \epsilon_{it} \tag{3.13}$$

where i indexes individuals and t indexes birth cohorts. The construction of the school-oriented and action-oriented trait-indices is described in the previous section. The test-score vector H can include test scores from mathematics, verbal, and elective subjects. Birth-cohort fixed effects δ_t are always included to facilitate pooled cross-sectional analysis. Earnings are calculated from the tax register as the sum of inflation-adjusted labor and entrepreneurial income averaged over age 35–38. An age interval is used to reduce measurement error and eliminate zeroes. The upper bound is chosen so that tax records exist for the last cohort (1979) in the last year of our main sample (2017).

Our main analysis sample includes male high-school graduates born in 1962–1979 for whom we have military test records. In the baseline estimation, all predictors are normalized to have zero mean and unit standard deviation within birth cohorts. High-school test scores are from a nationwide high-school exit exam taken around age 18. The military test is standardized also across cohorts and completed shortly after high school during basic training. All tests are graded in a double-blind procedure.

Table 3.2 presents the estimates of the β coefficients at different stages of saturation.[15] Column 1 shows that action-oriented and school-oriented traits have independent predictive power on earnings. The standardized coefficients of both measures are statistically significant and large in magnitude. The action-oriented trait has a lower earnings premium at 5.3 log points per standard deviation increase in the trait, while the premium for the school-oriented trait is almost twice as

[15] In Appendix C.3, Table C.2 shows the results when the school-oriented and action-oriented indices are replaced with the original personality measures. Table C.3 shows that the observed patterns are robust to using levels of earnings. Table C.4 shows that the results are not sensitive to truncating the long left tail in the log earnings distribution.

large, 9.6 log points.

Column 2 shows the coefficient for high-school mathematics score without controlling for traits. The earnings premium for mathematics is 16.4 log points per one standard deviation increase in the test score. Column 3 estimates the returns to the action-oriented trait, school-oriented trait, and mathematics in the same regression. Compared to column 1, the action-oriented trait's coefficient is almost double and, the school-oriented trait's coefficient is less than half. In contrast, the coefficient for mathematics barely moves.

Column 4 displays the estimates with additional high-school test score measures and IQ. The results reinforce the pattern observed in Column 3: The action-oriented trait coefficient further increases, and the coefficients for the school-oriented trait and mathematics decrease. The results imply that conditional on a comprehensive battery of standardized measurements around age 18, men who rank one standard deviation above the mean in the action-oriented trait earn over 11 log points more relative to the mean ranked men. This is comparable to men who score one standard deviation higher in the mathematics test and earn a 13 log point earnings premium. In this specification, only the test scores for elective subjects hold any non-negligible predictive power over mathematics scores and the action-oriented trait. For simplicity, in further analysis, we compare results from specifications 1 and 3.

In view of our model, estimates from Column 1 correspond to the marginal returns of the endowments (Expression 3.11). Conditional on the school-oriented trait, the returns to the action-oriented trait comprises of two opposite effects: the positive direct effect of having higher informal skills and the negative indirect effect due to decreased study effort. The positive sign of the action-oriented trait implies that the direct effect, on average, dominates the indirect effect. The results are also consistent with the prediction that the marginal returns to the school-oriented trait are always positive.

Empirically, the behavior of the action-oriented and school-oriented coefficients is best understood in the light of the results from Section 3.4.1. If academic performance is rewarded in the labor market, the school-oriented trait should have a positive premium in a regression without test score controls, because school-orientation is positively correlated with test-score performance. Additionally, school-oriented traits could improve earnings more directly if, for example, they foster noncognitive skills that are not particularly useful in studying but still improve job performance in some tasks. We do not find that the school-oriented trait is independently valued in the labor market, as evidenced by the small coefficient in the baseline specification. This is in contrast with studies that look at the correlation of personality traits with earnings without taking into account school performance (Almlund et al., 2011).

Our model abstracts from unobserved heterogeneity in students' preferences. However, an unobserved component is necessary to make sense of estimating traits and test scores in the same regression---otherwise there would be no variation in test scores conditional on traits. If the unobserved tastes are uncorrelated with ϵ_i (conditional on N and J), then the coefficient of mathematics corresponds to the returns r_H in our model. The coefficient of J, on the other hand, could be interpreted as the direct effect in Expression 3.11. The indirect effect is zero, because controlling for test scores holds socializing constant ($\partial s^*(N,J)/\partial J = 0$).

Note that while there is a mechanical aspect to the shrinking of the school-oriented trait due to its definition, is not mechanically driven to zero if there were to exist a direct channel of influence ($r_N > 0$). Additionally, IQ behaves in a very similar way to the school-oriented trait in the regressions. One interpretation is that the labor market rewards IQ and school-orientation because it enables success in tasks similar to test performance.

The two different estimates for the action-oriented trait can be understood in terms of timing. The estimate in Column 1 represents the effect of having a different endowment at the beginning of schooling. The estimate in Column 3, on the other hand, represents what the returns to informal skills would be if informal skills could be altered independently of educational capital, for example, after already completing education.

3.4.2.1 Returns to Skills within Occupations and Education

Why is the worse predictor of school performance such a strong predictor of earnings? In this section, we discuss the potential mechanisms that give rise to these premiums. We analyze selection into different education paths and occupations, and experience, career advancement, and job performance within occupations.

Table 3.3 presents results from the baseline regression (3.13) where fixed effects are progressively added for the level of education, occupation, and firm.[16] In column 2, the inclusion of occupation and education fixed effects shrinks the coefficient of school-oriented from 0.096 to 0.013, over an 85% decrease. Conversely, the coefficient of the action-oriented trait is almost unchanged from 0.53 to 0.51. This suggests that a large part of the premium for the personality that predicts school performance arises from sorting into profitable education paths and higher-paying occupations. In contrast, the premium for the action-oriented trait is less affected by sorting.

Are action-oriented men able to sort into higher-paying firms within the same occupation? Column 6 adds a firm fixed effect in addition to occupation and education fixed effects. This regression already explains 58% of the variation in adult earnings. While reducing the action-oriented trait's coefficient, its premium is still economically significant and almost twice as large as the coefficient for mathematics. The results imply that action-oriented (one standard deviation above the mean) employees earn 3.5 log points more relative to their colleagues with the same occupational and educational background even within the same firm.

Two caveats relate to these regressions. First, due to selection, labor market outcomes such as firm, education, and occupation are fundamentally 'bad controls' in an earnings regression. For example, Column 6 implies a comparison between men with average action-oriented trait and men with a high action-oriented trait, who are highly educated and working in a high paying occupation. There are likely to be unobservable reasons why these two different types of men would end up in a

[16] Some of these variables are included only for a subset of cohorts. This is reflected in the sample size of the regressions. The resolution of these variables consists of 66 categories for 'level × field of education' ('master's in humanities') and 80 harmonized occupational categories. The variables are recorded at the age of 35.

similar job and education. If the same unobservables influence earnings, it would bias the estimate for the returns to the action-oriented trait *even* if the action-oriented trait was a randomly assigned endowment. However, we find it warranted to draw attention to these suggestive results.

Second, the resolution of the fixed-effect variables matters for the size of the coefficients. Comparing within ever smaller groups almost necessarily decreases the coefficients by accounting for unobservable dissimilarities across larger groups. For this reason, we focus the attention to the relative magnitudes between the action-oriented and the school-oriented traits.

In summary, educational and occupational sorting appear to play an important role in the channel through which personality and academic achievement influence earnings. However, while still considerable, their role appears to be less important for determining the returns to the 'action-oriented' trait.

3.4.2.2 Occupational Sorting

In this section, we look at occupational sorting in more detail. For the analysis, we estimate Equation 3.13 with different response variables. We present results with and without mathematics included. Table 3.5 presents the results when mathematics is included. Column 8 shows that mathematics score and the action-oriented trait have the opposite impact on years of education. A one standard deviation increase in the mathematics scores predicts a 0.9 increase in years of schooling. Conversely, a one standard deviation increase in the action-oriented traits predicts a -0.3 decrease in years of schooling.

Consistent with higher educational attainment, high mathematics scorers work in professional occupations. Columns 1-6 of Table 3.5 use occupational indicators as the response variable. One standard deviation increase in the mathematics test score predicts a 13 percentage point increase in the probability of working in a professional occupation at age 35. Conversely, a similar increase in the action-oriented trait predicts a 6 percentage point decrease in the probability of working in a professional occupation.[17] On the other hand, a one standard deviation increases in either the mathematics test score or the action-oriented trait increase the probability of being a manager by 2 percentage points or 20%, taking into account the 10% baseline fraction of managers.[18]

Finally, Column 7 uses the average earnings in the individual's occupation as the response variable. The results show that a one standard deviation increase in mathematics test scores is associated with being employed in an occupation with 12 log points higher earnings. Higher action-oriented traits do not predict working in a high paying occupation. In other words, despite earning more themselves, action-oriented men do not work in particularly high-paying occupations. This is consistent with the results from occupational sorting (Columns 1-6). Unlike the mathematics score, the action-oriented trait shows no clear pattern predicting sorting away from low-skill occupations.

How do men spend their years between high school graduation and age 38? The response

[17] Relative to the baseline of 42%, these numbers correspond to a 31% increase and a 14% decrease in the likelihood of working in a high-level professional occupation.

[18] Our definition of managers excludes small-business owners who perform employee-type work in the firm, such as, owners of small trucking firms.

variables in Table 3.6 are cumulative years spent in the given activity from age 18 to 38. Columns in Table 3.6 represent the exhaustive and mutually exclusive list of principal activities recorded by Statistics Finland yearly for each person. By construction, each row sums to zero. The results indicate that a mathematics test score one standard deviation above the mean is associated with 0.43 years of study, and 0.46 years less nonemployment. Conversely, a one standard deviation increase in the action-oriented trait is associated with 0.48 fewer years of studying, 0.66 more years of work experience, and 0.15 fewer years of nonemployment relative to the average individual. In other words, for high mathematics men, the time spent studying is fully offset by reduced nonemployment instead of reduced work experience, whereas for the action-oriented men, the offset for the increase in work experience comes from both less time in nonemployment *and* less studying.

In summary, men with a personality that predicts low school performance start their careers earlier, accumulate more work experience by avoiding nonemployment and skipping education, and are more likely to end up in managerial positions relative to their more average peers. They enjoy an earnings premium even without placing in particularly high-paying occupations. Conversely, school achievers manage to educate themselves without compromising work experience. Relative to the average high school graduate, they are more likely to be employed in high-paying professional and managerial occupations. Together, we interpret this as suggestive evidence that action-oriented traits help workers by improving the gradient of their career progress. In contrast, school-oriented traits and mathematics ability help workers to start higher up on the ladder.

3.4.3 Time Trends in Skill Premiums and Skill Specialization

In the cross-sectional analysis, we show that both the school-oriented and action-oriented traits have considerable earnings premiums. Figure 3-3 shows how these premiums have changed over time by estimating Equation 3.13 for each birth cohort separately (omitting the cohort fixed effect). The results show a striking reversal in the magnitude of the premiums over the 16-year period. The premium for the action-oriented trait has increased from virtually zero to almost 9 percentage points per standard deviation. The school-oriented trait has decreased from 0.13 to 0.8 over the same period. Appendix C.3 Figure C-2 shows the corresponding figure with mathematics included in the regression.

Our result is consistent with Deming (2017) who finds an increase in the returns to a proxy of sociability between the NLSY79 and NLSY97 cohorts and with Edin et al. (2021) who find an increase in the returns to noncognitive skills in Sweden. However, we show that noncognitive skills are inadequately described by a single dimension, as demonstrated by the opposite trends in the returns to the action-oriented and school-oriented traits. Furthermore, the action-oriented trait appears to measure more than just teamwork-based skills as in Deming (2017). We observe that the increasing trend for the action-oriented trait returns is driven by all three of its components, sociability, activity, and masculinity, not only sociability.[19]

Deming (2017) offers a demand-side explanation of the growing importance of social skills:

[19]Results are available by request.

changing job requirements increase the demand for social skills. Deming shows how the equilibrium employment shares have changed to favor socially intensive jobs at the expense of mathematics-intensive jobs. In the same spirit, but on the supply side, we use tercile cutoffs within cohorts to group men into four bundles: 'school-specialized' (high mathematics, low action-oriented), 'action-specialized' (high action-oriented, low math), both high (high math, high action-oriented) and both low (low math, low action-oriented).

Skill specialization implies inverse bundling of action-oriented personality traits and mathematics skills; we should observe relatively fewer men who rank high in both. Figure 3-4 shows the evolution of the relative shares of each bundle over time. Each cohort is divided into mutually exclusive groups (bundles) along the tercile cutoffs in their mathematics score and action-oriented trait. Each line represents the evolution of the size of the bundle. Figure 3-4 shows clearly that skill specialization has increased over time. The largest divergence happens in the latter half of the period (individuals for which labor-market earnings are measured after 2005). Before that, the relative proportions of the bundles are roughly equal. In the last cohort, there are 5 percentage points more inverse bundles (action-specialized and school-specialized) relative to the 'pooling' bundles. From the baseline of 11% share each, this corresponds to 10% increase in the inverse bundles and a 10% decrease in the pooling bundles.

What is driving the increasing separation of school performance and action-oriented traits? If specialization is an equilibrating reaction from the supply side to the increased demand for social skills in the labor market, all students should increase their informal-skill investment. However, students with higher marginal benefits should do so relatively more. In our model, if the marginal cost of studying is increasing in the comparative advantage to the action-oriented trait, we should see the largest time-reallocation towards informal skills for those who already have a comparative advantage in action-oriented traits. Likewise, earnings gains should be largest for 'action-specialized' and smallest for 'school-specialized.' Earnings gains for the separating bundles should fall somewhere in between.

Figure 3-5 shows the evolution of earnings for each group relative to the cohort born in 1963.[20] Changes in earnings have been uneven across the bundles. Those with high 'action-oriented' traits and low mathematics skills ('action-specialized') had a 20% increase earnings. The earnings of 'school-specialized' improved by 10% in the same period. The changes in earnings of high-math/high-action-oriented individuals and low-math/low-action-oriented individuals place between the inverse bundles.

In summary, our results are consistent with a supply-side response to the increased returns of social skills. Simultaneously, we emphasize that while the novel descriptive trends are robust, a full explanation for the trends requires further research.

[20]We omit the 1962 cohort because it has substantially fewer observations.

3.4.4 The Reading Penalty

This section relates our results to the 'reading penalty' feature found in several US longitudinal data sets. Sanders (2015) and Altonji et al. (2016) show that in a wage regression that includes both mathematics and verbal test scores as regressors, the verbal score has a negative coefficient. Sanders (2015) demonstrates that this is a robust feature of five commonly used US longitudinal data sets.[21] That study controls for occupational and educational sorting and crude measures of personality but still finds a negative partial correlation between the verbal test scores and wages.

Column 1 of Table 3.7 shows the results of estimating Equation 3.13 in our data with only high-school test scores as regressors. We also find a small negative coefficient of -0.6 log points for the verbal test. Similarly, controlling for education and occupation in Columns 2 and 3 only serve to reduce the math and electives coefficients, but not the one for verbal scores. We also conclude that differential occupational sorting is not the source of the reading penalty.

Column 4 adds personality and IQ controls to the regression. In this specification, verbal scores have a slightly positive coefficient of 0.5 log points. While not conclusive, the evidence supports the hypothesis that inverse bundling of personality and verbal skills contributes to the observed reading penalty, and that the returns to verbal skills in the labor market are low. In view of our framework, action-oriented students invest less in verbal skills. If the returns to verbal skills are particularly low, the only ones investing in them are students who have a high comparative advantage in the school-oriented trait.[22]

3.5 Conclusion

This paper analyzes how do different dimensions of personality predict school vs. labor-market performance, and how the labor-market value of these traits has changed over time. It uses data that includes multidimensional personality and cognitive ability measures, education, and labor-market records for 79% of Finnish men.

We demonstrate that to understand the role of noncognitive skills in the labor market it is essential to consider the multidimensional nature of skills. The key reason is that different dimensions of noncognitive skills appear to have opposite effects in human capital production relative to the labor market. At its core, the paper considers the distinction between the school versus the labor-market. We find that high achievers in school lack, on average, at least in some dimensions of economically valuable personality traits. Conversely, low-achieving students tend to have some redeeming qualities that compensate in the labor market for their lack of academic success.

We formalize this idea using a model of multidimensional skill specialization. Variation in initial personality endowments generates differences in comparative advantage that leads to specialization in 'school-orientation' and 'action-orientation.' We explore the empirical implications of this model

[21] NLSY79, NLSY97, NELS88, ELS88 and Baccalaureate and Beyond (Sanders, 2015).

[22] In an extension to our framework, productivity for mathematics and verbal skills could depend on different kinds of traits which we measure imperfectly.

on educational and occupational sorting and careers. Using the model to structure our analysis, we provide four new empirical facts.

First, a particular subset of personality traits predicts high educational achievement, but another critical subset of personality traits predicts low achievement. This pattern follows the common stereotypes: men, who score high in sociability, activity-energy, and masculinity, tend to perform worse in standardized school tests. Achievement striving, dutifulness, and deliberation predict good school performance.

Second, the traits that predict low school achievement still predict labor market success. Conditional on test scores, one standard deviation increase in these traits predicts a 10 log point increase in earnings at age 35. The corresponding returns to mathematics in the same regression is 16 log points. In contrast, the traits that predict higher school achievement are not independently valued in the labor market.

Third, the economic returns to traits that predict low school performance have rapidly increased over the past two decades. Men with high sociability, activity-energy, and masculinity, but with low math skills, experienced the highest earnings gains. Returns to traits that predict high school achievement have declined, and cognitive skills returns have been stable.

Fourth, skill specialization has increased over the past two decades: men have become more likely to possess either good formal or informal skills and are less likely to have both.

Main Figures and Tables

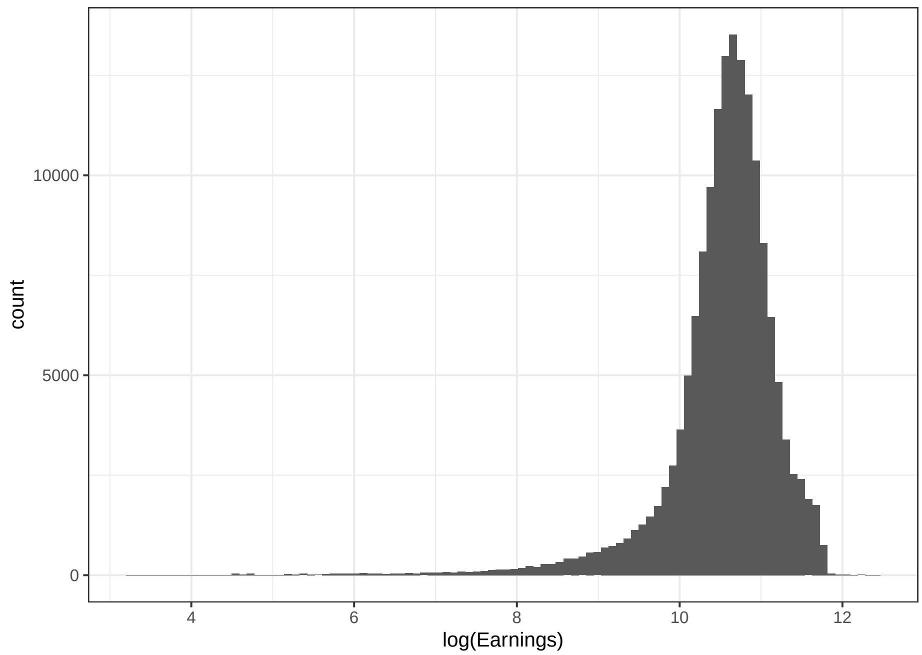

Figure 3-1: Distribution of Log Earnings.

Notes: The histogram shows the distribution of log earnings in the main sample of male high school graduates with military test scores. In the sample, n = 158,000, mean = 10.5, SD = 0.72.

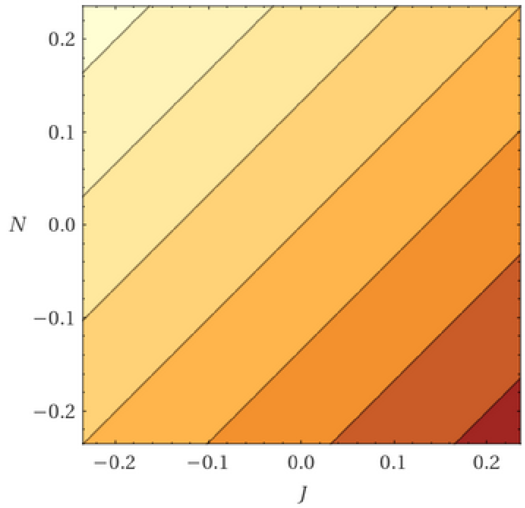

Figure 3-2: Comparative Advantage.

Notes: Each line represents an isoquant in a plane where J and N are in the x and y axis and $s^*(N, J)$ is in the z axis. Darker shades indicate higher values of z. Functional form choices are $a(N, J) = N$, $b(N, J) = J$, $C(1 - s, N, J) = (J + (1 - s))^2$.

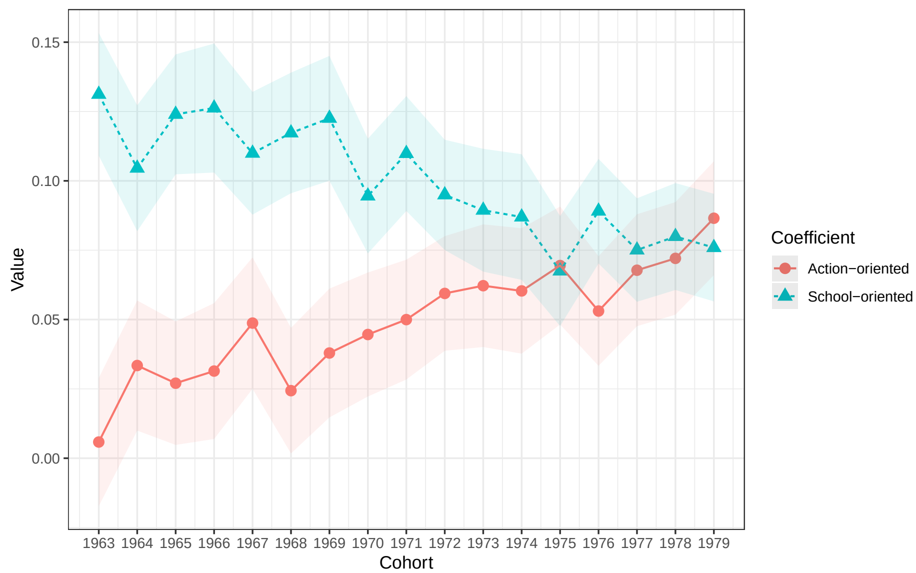

Figure 3-3: Time Trends in Returns to Traits.

Notes: Each point in the figure corresponds to a regression coefficient from estimating Equation 3.13 separately for each cohort, with log earnings as the outcome and person as the unit of observation. The right-hand-side variables include only the action-oriented and school-oriented traits. The action-oriented trait is a composite of Sociability, Activity, and Masculinity. The school-oriented trait is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. All covariates are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are recorded by the tax authorities and measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported as the shaded area.

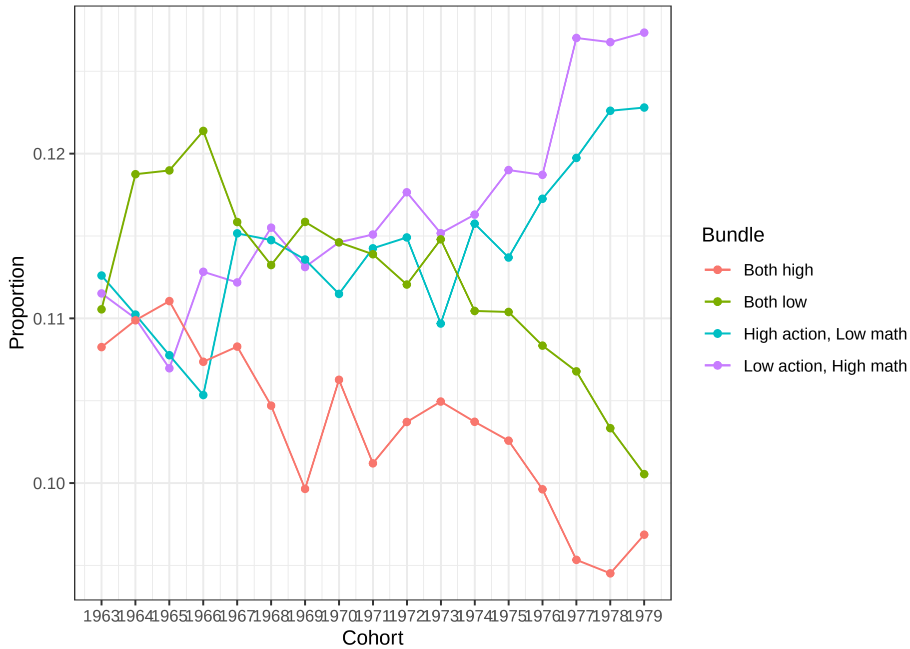

Figure 3-4: Time Trends within Bundles.

Notes: Each point corresponds to the proportion of persons belonging to the indicated group (bundle) within that cohort. The 'High action, Low math' bundle includes persons who belong to the top tercile in the action-oriented trait and the bottom tercile in the math score. The 'Low action, High math' bundle includes persons who belong to the bottom tercile in the action-oriented trait and the top tercile in the math score. The 'Both high' bundle includes persons who score in the top tercile in both dimensions and vice versa for the 'Both low' bundle.

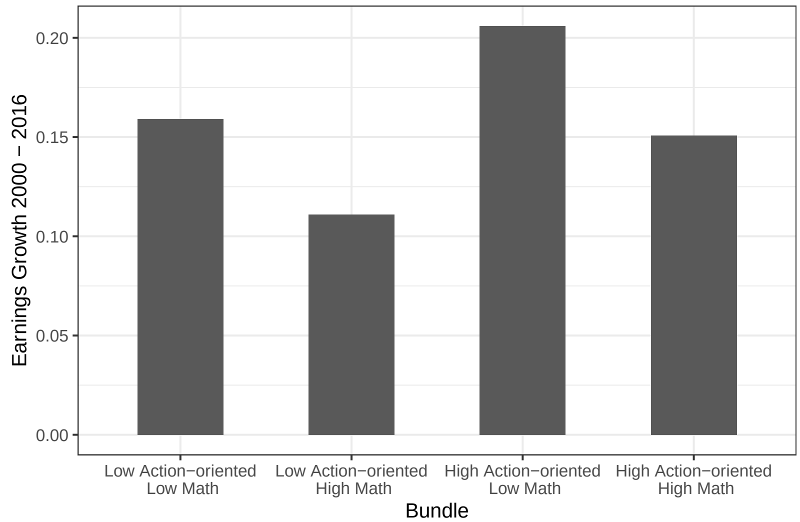

Figure 3-5: Earnings Change within Bundles.

Notes: Each bars corresponds to the change median earnings of that bundle relative to the base year 1963. The 'High action, Low math' bundle includes persons who belong to the top tercile in the action-oriented trait and the bottom tercile in the math score. The 'Low action, High Math' bundle includes persons who belong to the bottom tercile in the action-oriented trait and the top tercile in the math score. The 'High action, High math' bundle includes persons who score in the top tercile in both dimensions and vice versa for the 'Low action, Low math' bundle.

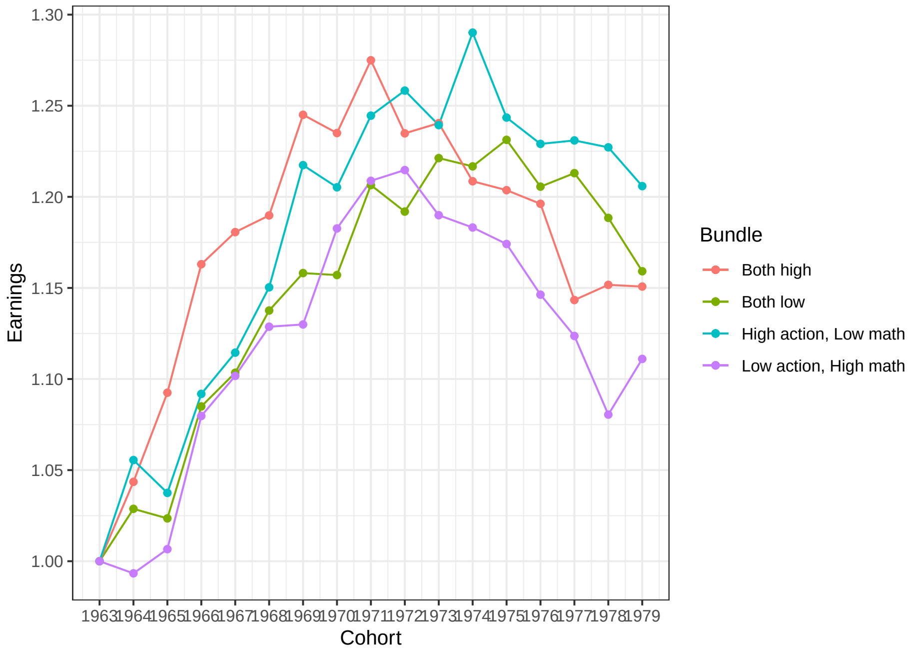

Figure 3-6: Earnings Change within Bundles.

Notes: Each point corresponds to the median earnings of that bundle relative to the base year 1962. The 'High action, Low math' bundle includes persons who belong to the top tercile in the action-oriented trait and the bottom tercile in the math score. The 'Low action, High math' bundle includes persons who belong to the bottom tercile in the action-oriented trait and the top tercile in the math score. The 'Both high' bundle includes persons who score in the top tercile in both dimensions and vice versa for the 'Both low' bundle.

Table 3.1: Personality and Academic Performance.

	High School Test Scores						
	Math	Verbal	Electives	HS GPA	In HS sample	9th grade GPA	Years of Education
Sociability	−0.258	−0.162	−0.219	−0.220	−0.054	−0.202	−0.335
	(0.004)	(0.004)	(0.004)	(0.004)	(0.001)	(0.004)	(0.006)
Activity	−0.122	−0.189	−0.145	−0.215	−0.099	−0.186	−0.448
	(0.004)	(0.004)	(0.004)	(0.004)	(0.001)	(0.004)	(0.005)
Masculinity	−0.032	−0.147	−0.126	−0.161	−0.059	−0.134	−0.222
	(0.003)	(0.003)	(0.003)	(0.002)	(0.001)	(0.003)	(0.004)
Deliberation	0.113	0.081	0.102	0.085	0.005	0.055	0.197
	(0.003)	(0.003)	(0.003)	(0.003)	(0.001)	(0.003)	(0.004)
Dutifulness	0.011	0.083	0.068	0.064	0.059	0.164	0.232
	(0.004)	(0.004)	(0.004)	(0.004)	(0.001)	(0.004)	(0.005)
Achievement Aim	0.190	0.168	0.198	0.224	0.111	0.301	0.636
	(0.004)	(0.004)	(0.004)	(0.004)	(0.001)	(0.003)	(0.005)
Confidence	0.263	0.151	0.181	0.219	0.097	0.249	0.508
	(0.004)	(0.004)	(0.004)	(0.004)	(0.001)	(0.004)	(0.005)
Leadership	0.070	0.111	0.142	0.121	0.081	0.091	0.257
	(0.005)	(0.005)	(0.005)	(0.005)	(0.001)	(0.005)	(0.006)
Y mean	0.000	0.000	0.000	0.000	0.360	0.000	12.890
Cohort FE	yes	yes	yes	yes	yes	yes	yes
Adj. R^2	0.090	0.095	0.111	0.128	0.206	0.249	0.190
Observations	157129	157129	150610	162605	459357	119902	457529

Notes: Each column reports the OLS regression results from Equation 3.12. The column name indicates the outcome. The unit of observation is the person. The standardized high school (HS) tests are administered by the Matriculation Examination Board before military service. Personality traits are measured by the Finnish Defence Force after high school. Test scores and personality traits are normalized to have mean 0 and standard deviation 1 within cohorts. All models control for the birth year (cohort) fixed effects. Robust standard errors are reported in parentheses.

Table 3.2: Returns to Skills.

	Dependent variable: log earnings			
	(1)	(2)	(3)	(4)
Action-oriented	0.053		0.099	0.112
	(0.003)		(0.003)	(0.003)
School-oriented	0.096		0.036	0.018
	(0.003)		(0.003)	(0.003)
Math		0.164	0.160	0.130
		(0.002)	(0.002)	(0.002)
IQ				0.012
				(0.002)
Verbal				0.005
				(0.002)
Electives				0.052
				(0.002)
Outcome mean	10.520	10.520	10.520	10.520
Cohort FE	yes	yes	yes	yes
Adj. R^2	0.048	0.063	0.093	0.098
Observations	157743	157891	157129	156843

Notes: Each column reports the OLS regression results from Equation 3.13, with log earnings as the outcome. The unit of observation is the person. 'Action-oriented' is a composite of Sociability, Activity, and Masculinity. 'School-oriented' is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. Test scores and traits are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported in parentheses.

Table 3.3: Returns to Skills within Occupation and Education.

	Baseline			With math control		
	(1)	(2)	(3)	(4)	(5)	(6)
Action-oriented	0.086	0.051	0.035	0.097	0.057	0.040
	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)
School-oriented	0.022	0.013	0.011	0.009	0.006	0.007
	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)
Math				0.066	0.035	0.028
				(0.002)	(0.002)	(0.002)
Outcome mean	10.520	10.520	10.790	10.520	10.520	10.790
Cohort FE	yes	yes	yes	yes	yes	yes
Education FE	yes	yes	yes	yes	yes	yes
Occupation FE	no	yes	yes	no	yes	yes
Firm FE	no	no	yes	no	no	yes
Adj. R^2	0.171	0.328	0.575	0.177	0.331	0.577
Num. obs.	157743	100472	61224	157129	100003	60940

Notes: Each column reports the OLS regression results from Equation 3.13, with log earnings as the outcome. All models control for the birth year (cohort) and additional fixed effects as indicated. Sample size varies when variables are available only for a subset of cohorts. When firm FE is included, public sector employees are excluded. The unit of observation is the person. The standardized high school (HS) tests are administered by the Matriculation Examination Board before military service. Personality traits are measured by the Finnish Defence Force after high school. The action-oriented trait is a composite of Sociability, Activity, and Masculinity. The school-oriented trait is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. Test scores and traits are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are recorded by the tax authorities and measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported in parentheses.

Table 3.4: Occupational Sorting.

	Occupations							
	Managers	Professionals	Technical/Clerical	Service/Sales	Production	Other	Mean Earnings	Years of Educ.
Action-oriented	0.014	−0.099	0.027	0.018	0.035	0.005	−0.030	−0.796
	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.005)
School-oriented	0.027	0.111	−0.038	−0.019	−0.066	−0.016	0.088	1.443
	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.005)
Outcome mean	0.100	0.420	0.290	0.060	0.100	0.030	10.680	12.890
Cohort FE	yes	yes	yes	yes	yes	yes	yes	yes
Adj. R^2	0.024	0.028	0.004	0.009	0.028	0.006	0.062	0.188
Observations	100472	100472	100472	100472	100472	100472	80186	457529

Notes: Each column reports the OLS regression results from Equation 3.13. The column name indicates the outcome. Each outcome variable is an indicator of working in the given occupation at age 35 (at the end of the calendar year). The outcome in the last column measures the average earnings of all other men employed in the same occupation at age 35. 'Action-oriented' is a composite of Sociability, Activity, and Masculinity. 'School-oriented' is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. Test scores and traits are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported in parentheses.

Table 3.5: Occupational Sorting with Mathematics.

	Occupations							
	Managers	Professionals	Technical/Clerical	Service/Sales	Production	Other	Mean Earnings	Years of Educ.
Action-oriented	0.020	−0.060	0.011	0.006	0.021	0.002	0.007	−0.297
	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.008)
School-oriented	0.019	0.063	−0.018	−0.003	−0.049	−0.012	0.043	0.568
	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.009)
Math	0.022	0.129	−0.053	−0.042	−0.046	−0.010	0.120	0.929
	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.006)
Outcome mean	0.100	0.420	0.290	0.060	0.100	0.030	10.680	14.740
Cohort FE	yes	yes	yes	yes	yes	yes	yes	yes
Adj. R^2	0.029	0.091	0.016	0.035	0.050	0.009	0.215	0.211
Observations	100003	100003	100003	100003	100003	100003	79804	156016

Notes: Each column reports the OLS regression results from Equation 3.13. The column name indicates the outcome. Each outcome variable is an indicator of working in the given occupation at age 35 (at the end of the calendar year). The outcome in the last column measures the average earnings of all other men employed in the same occupation at age 35. 'Action-oriented' is a composite of Sociability, Activity, and Masculinity. 'School-oriented' is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. Test scores and traits are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported in parentheses.

Table 3.6: Cumulative activity, Ages 18-38.

	Baseline				With math control			
	Work Experience	Study	Nonemployment	Other	Work Experience	Study	Nonemployment	Other
Action-oriented	0.647	−0.607	−0.015	−0.024	0.662	−0.477	−0.150	−0.036
	(0.015)	(0.012)	(0.008)	(0.005)	(0.016)	(0.012)	(0.008)	(0.005)
School-oriented	−0.139	0.408	−0.314	0.045	−0.162	0.247	−0.146	0.060
	(0.015)	(0.011)	(0.008)	(0.004)	(0.016)	(0.012)	(0.008)	(0.005)
Math					0.062	0.432	−0.455	−0.038
					(0.011)	(0.008)	(0.006)	(0.003)
Outcome mean	14.360	4.510	1.050	1.080	14.360	4.510	1.050	1.080
Cohort FE	yes	yes	yes	yes	yes	yes	yes	yes
Adj. R^2	0.039	0.033	0.041	0.021	0.039	0.060	0.103	0.022
Observations	98138	98138	98138	98138	97658	97658	97658	97658

Notes: Each column reports the OLS regression results from Equation 3.13. The column name indicates the outcome. The outcome variable is measured in years. The unit of observation is the person. The standardized high school (HS) tests are administered by the Matriculation Examination Board before military service. Personality traits are measured by the Finnish Defence Force after high school. The action-oriented trait is a composite of Sociability, Activity, and Masculinity. The school-oriented trait is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. Test scores and traits are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are recorded by the tax authorities and measured by averaging total labor and entrepreneurial income earned at age 35-38. All models control for the birth year (cohort) fixed effects. Robust standard errors are reported in parentheses.

Table 3.7: Reading Penalty.

	Earnings	Earnings	Earnings	Earnings
Math	0.138	0.056	0.027	0.130
	(0.002)	(0.002)	(0.002)	(0.002)
Verbal	−0.006	−0.010	−0.004	0.005
	(0.002)	(0.002)	(0.002)	(0.002)
Electives	0.056	0.024	0.018	0.052
	(0.002)	(0.002)	(0.002)	(0.002)
Action-oriented				0.112
				(0.003)
School-oriented				0.018
				(0.003)
IQ				0.012
				(0.002)
Outcome mean	10.520	10.520	10.520	10.520
Cohort FE	yes	yes	yes	yes
Education FE	no	yes	yes	no
Occupation FE	no	no	yes	no
Adj. R^2	0.068	0.158	0.319	0.098
Observations	157605	157605	100031	156843

Notes: Each column reports the OLS regression results from Equation 3.13, with log earnings as the outcome. All models control for the birth year (cohort) and additional fixed effects as indicated. The unit of observation is the person. The standardized high school (HS) tests are administered by the Matriculation Examination Board before military service. Personality traits are measured by the Finnish Defence Force after high school. The action-oriented trait is a composite of Sociability, Activity, and Masculinity. The school-oriented trait is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. Test scores and traits are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are recorded by the tax authorities and measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported in parentheses.

Chapter 4

The Surprising Intergenerational Effects of Manufacturing Decline

4.1 Introduction

This paper is about children of the left-behind places of America—the children of crisis. It asks what happens to children in the many declining manufacturing towns and cities of the US. The main focus is whether the local decline in manufacturing employment has increased—or decreased—the high-school dropout rate. The paper also explores the consequences of manufacturing decline on educational mobility—that is, the chances that a child born to poor parents enrolls in a college, and the factors that characterize the places with the largest effects on children.

The American middle class has declined across the country, affecting places from Detroit to Boston, from Middletown, Ohio to Washington, DC. The main causes of this—technology and trade—have eliminated a large part of US manufacturing jobs, and plausibly continue to do so (Acemoglu et al., 2016). The effects are most visible at the geographical level: some places have been left-behind while some places prosper. The haves and have-nots live in different places (Moretti, 2012; Florida, 2017). This is well documented: geographically uneven manufacturing decline and shrinking middle incomes are the key factors in America's deepening divide between rich and poor (Autor et al., 2015).

The previous research on globalization, technology, and inequality focuses primarily on adult males' labor market prospects. While this is undoubtedly important, the long-term effects—the future of work—depend on children, the next generation. This aspect of labor market adjustment—children—has received surprisingly little attention in the literature. But it could be the most important margin.

To study this, the paper uses county-level data on the employment structure and children's educational outcomes from the U.S. 1990–2010. To establish causal inference, the paper uses variations in trade exposure from China following its entry to the World Trade Organization (WTO) as an instrument for local manufacturing decline in the US. The instrument is computed from detailed product-level trade data from the UN Comtrade database. To explore the local factors correlated with the effects, the paper uses a large set of data on community characteristics, from segregation to educational resources. The idea of the empirical setup is that, conditional on the instrumental variables strategy, otherwise similar places faced different levels of manufacturing decline. This identifies the effects on children.

The literature on manufacturing decline of the 21st century paints a bleak picture. In places that have been hit the hardest, workers—especially adult men—have been slow to adjust (Autor et al., 2014; Yagan, 2017). These places are characterized by job losses, lower employment and wages, and increased applications for social assistance (Autor et al., 2013; Balsvik et al., 2013). Contemporary evidence also suggests that manufacturing decline is a source of social distress. When factory jobs vanish, men become less desirable partners and divorces more common (Autor et al., 2017). Violent- and property-crime tend to increase (Pierce and Schott, 2016b; Deiana, 2016; Feler and Senses, 2017). In places that experienced trade-induced manufacturing decline, children become more likely to be raised in poor single-headed households, making childhood poverty more prevalent (Autor et al., 2017). Based on this evidence, it would be reasonable to conjecture that manufacturing decline could make teenagers more prone to drop out of high-school and direct them away from college.

This paper finds the opposite. In places where manufacturing has declined, children drop less out of high-school. The relationship appears to be causal: comparing places within the same US region, with similar initial share of workers employed in manufacturing, and with similar demographic characteristics; those places that saw manufacturing decline because they were historically specialized in the particular industries that China started to export in 2001, saw sizable decreases in high-school dropout rates—compared to the otherwise similar places that were not exposed to competition with China. This paper also finds that when manufacturing employment declines, chances that poor children enroll in college increase. The causal evidence on the second observation is less conclusive but it is consistent with the first finding.

The paper also analyzes the local characteristics that could mediate, mitigate, or amplify the effects. To do so, it estimates interactions between manufacturing decline and a large set of factors that have been discussed in the sociology and economics literature, such as segregation and inequality. In contrast with the literature on the determinants of upward income mobility, I find that the effects are larger in areas with higher segregation and with larger African American populations. Local educational resources, such as school spending or student-teacher ratios show no significant correlations with the size of the effect. If anything, their predictive effect is negative. These are new and puzzling findings.

The main results are consistent with the idea that the manufacturing decline increased returns and decreased opportunity costs of education, and with sociological accounts linking working-class environment and children's education. In the classical Becker (1964) model of human capital investment, the decision-maker—in this case a teenager—compares the marginal costs and benefits of education. Complementary evidence by Autor et al. (2013) shows that trade-induced US manufacturing decline reduced the wages for individuals with low levels of education, compared to those with more, plausibly increasing the relative benefits of schooling. On the opportunity cost side, a reduction in available manufacturing jobs may have reduced the outside options for high-school dropouts.

From sociological perspective, Willis (1977), in the landmark research "*Learning to Labor: How Working Class Kids Get Working Class Jobs*", highlights how children inherit occupations and class from their parents and community. In working-class communities, Willis (1977) notes, counter-school culture of resistance and opposition to academia are prevalent. But possibly a decline in working-class jobs, as in this paper, could lead to a decline in working-class culture. Following Willis (1977)'s argument, this could lead to an increase in children's education. Willis' theory could also help reconcile the interaction effects between local segregation and manufacturing decline: more segregated places could be the ones supporting stronger and more uniform working-class culture. When factory jobs vanish, the culture fostering high-school dropout behavior could dissolve, especially so in segregated places where the local culture may have been stronger.

In contrast to the group-level analysis of this study, a body of literature studies the individual-level effects of parental job loss. Most of it finds negative effects. For example, Oreopoulos et al. (2008) find that children whose fathers were displaced face long-lasting effects into adulthood: lower earnings, higher social assistance, and lower college attendance. Other longitudinal studies find that parental job loss decreases school grades (Rege et al., 2011) and increases grade repetition (Stevens and Schaller, 2011).[1] But these opposite results do not need to be contradictory. Those children whose parent lost a job tend to be negatively affected, but—at the local level—the other children could primarily respond to the changed incentives and local environment—returns to education and the opportunity cost of it—while avoiding the cost of job loss in the family.

When factories closed in the US, some new factories opened in the developing world. In line with the results of this study, Atkin (2016) finds that local factory *openings* in export-manufacturing industries lead to *higher* school dropout rates in Mexico. Young people dropped out of school to work in manufacturing. This is a mirror image to what appears to have happened in the US. The effect is reasonably identified: Atkin (2016) uses the variation in the timing of factory openings across commuting zones in Mexico during a period of major trade reforms 1986–2000. Atkin (2016) argues that the effects are driven by the increased opportunity cost of schooling.[2]

This study's results are also consistent with the available local evidence from the US. Using historical data, Goldin and Katz (2000), show that industrialization slowed the growth in high

[1]Much fewer and less strongly identified studies focus on the community-level effects of job losses. A series of papers by Ananat et al. (2011) and Ananat et al. (2017) explore this aspect by comparing U.S. states. In line with individual-level effects, they find large negative effects on student achievement and college mobility from state-level job losses. The correlations they document at the state level, however, may not need to be causal. Another interpretation is that they focus on different type of variations in job losses.

[2]Similarly, Shah and Steinberg (2017) observe that in India children dropped out of school into productive work when rainfall was higher. In their setting, the opportunity cost of schooling, even for fairly young children, appears to have been an important factor in determining overall human capital investment. Munshi and Rosenzweig (2006); Shastry (2012); Jensen (2012) and Oster and Steinberg (2013) provide complementary evidence on the arrival of high-skill service jobs in India.

school attendance in the early 20th century United States. Focusing on the Appalachian coal boom and bust of the 1970s and 1980s, Black et al. (2005) find that the boom lead to increases in school dropout rates and the bust decreased them.

This analysis on the intergenerational effects of manufacturing decline builds on the work of Autor et al. (2013)—and related studies by Acemoglu et al. (2016), Pierce and Schott (2016a), and Bloom et al. (2016)—by using the rapid expansion of China's exports in manufacturing goods for empirical identification. Among other results, Autor et al. (2013) confirm the classical prediction Heckscher–Ohlin model of international trade: there are winners and losers from the geography of globalization. This research paper expands their work to include analysis on the consequences of the manufacturing decline more generally, a dimension they do not consider, and by characterizing the intergenerational effects on children's education. In short, Autor et al. (2013) focus on the causes and consequences of the 1990–2007 US manufacturing decline on adult men. This paper looks at the intergenerational effects of it.

The contribution of this paper is empirical: it answers the question of how children have been affected by the rapid manufacturing decline of the 2000's in the US. To date, this question has had no attention, or an answer in the literature. This research matters because the future of work critically depends on the labor market prospects of the next generation. The paper provides new evidence on how people and communities adjust to the structural transformation of work.

The article is organized as follows. Section 4.2 describes the data set and the empirical methodology. Section 4.3 reports the primary ordinary least squares (OLS) and two-stage least squares (2SLS) estimates of the impact of manufacturing decline shocks on high-school dropout rates and college mobility. Section 4.4 explores the robustness of the main results through several tools. Section 4.5 takes the analysis further and explores interactions between the effects of children and observable characteristics of commuting zones. Section 4.6 discusses the findings, provides interpretations, and connects the results to earlier empirical literature. Section 4.7 concludes.

4.2 Empirical Approach

4.2.1 Local Labor Markets

The unit of the analysis is regional economies—the local labor markets of the United States. The idea of the geographical analysis is that strong regional variations in the industry specialization make different places differentially exposed to shocks in manufacturing employment. Decline in manufacturing has varied by region and over time, not at the individual level, making local economies a natural observation unit. The operational geographical units are 722 commuting zones (CZ) developed by Tolbert and Sizer (1996). They approximate the areas where the population of interest works.[3] The CZ:s cover all metropolitan and nonmetropolitan areas, both urban and rural, of the mainland United States. The CZ:s are based on economic geography rather than

[3] Tolbert and Sizer (1996) measure commuting ties between US counties and define commuting zones as collections of counties with strong commuting ties between them.

administrative borders, are time-consistent, provide more granular measurement than state-level analysis, and can be matched to various official statistics (Autor and Dorn, 2013). Table 4.2 summarizes descriptive statistics for the CZs.[4] Figure 4-3 displays a map of the CZs, with the key variables of this study.[5]

4.2.2 Manufacturing Decline

A. Descriptive Data

The main data source on the US employment structure is the County Business Patterns (CBP) from 1991 to 2011 provided by the US Census Bureau. The CBP provides annual data on employment and payroll by county and industry. The data cover all US private employment, excluding most government employees, agricultural workers, self-employment, private household employment, and railroad workers.[6]

To complement the employment statistics, the paper uses population data from the Census Population Estimates. It provides data on the total and working-age (ages 15–64) US population at the county level. The county level data are mapped to CZs using the matching strategy detailed in Dorn (2009).

The main explanatory variable is the (annualized) decadal change in the share of manufacturing employment E_i^{MF} within total employment E_i^{TOT} in a CZ i:

$$\Delta MF_{it} = \Delta \left(\frac{E_{it}^{MF}}{E_{it}^{TOT}} \right). \tag{4.1}$$

Figure 4-1 describes the evolution of the US manufacturing employment based on CBP data. The US manufacturing employment was approximately constant in 1991–2000, but declined rapidly by 33.3 percent in 2000–2011.[7] Manufacturing's share of total employment was 19.1 percent in 1991 and fell to 10.4 percent in 2011. The rate of change in the manufacturing employment also had large variations between CZ:s and over time, as shown in the map of Figure 4-3 and in Table 4.3. While the manufacturing share of employment decreased on average in the US over 1991–2011, some places saw even increases in it. Table 4.2 summarizes descriptive statistics of manufacturing-to-total employment ratios, as well as employment-to-population ratios and the population size of the CZ—key baseline control variables in the estimation.

[4] An average CZ had a population of 350,000 in 1991. The largest commuting zone, New York, NY, had a population of 10.4 million and the smallest had 1,311.

[5] The object of interest in this study is the childhood environment, and therefore this analysis treats CZ:s as the observation units of interest without weighting them.

[6] For confidentiality reasons CBP reports employment by industry as an interval. I compute employment in these cases using the fixed-point imputation strategy developed by Autor et al. (2013).

[7] US manufacturing employment was 17.0 million in 1991, 17.1 million in 2000, 13.9 million in 2007, and 11.4 million in 2011, according to CBP data.

B. IV Strategy

To identify plausibly exogenous variations in manufacturing decline, I use an instrumental variables strategy (IV) based on the local industry exposure to China's imports.[8] Between 1990 and 2011, the share of US manufacturing imports from China increased over four-fold, from 4.5 percent to 23.1 percent (Fig 4-2).[9] This increase coincides with a sharp drop in the US manufacturing employment after 2000 (Fig 4-1).

The general idea is that China's entry to the world market is close to an exogenous shock to US manufacturing labor demand (Autor et al., 2013). The increase in China's exports to the US originates from China, not the US. It was sparked by China's large economic reforms in 1980–2000, and made possible by two sudden policy changes in 2001: China's accession to the World Trade Union (WTO) and a change in a US trade policy that eliminated potential tariff increases on Chinese imports (Pierce and Schott, 2016a; Hanson, 2012; Naughton, 2006). China's exports to the US were almost exclusively in manufacturing goods. This translated to a negative shock to US manufacturing labor demand in the 1999–2011 (Autor et al., 2013).[10]

The particular implementation of the IV strategy originates from the approach of Autor et al. (2013) using local labor market variations in the US industry exposure to Chinese import competition. The measure of exposure to China's imports leverages the fact that commuting zones vary in their distribution of industrial employment, making some commuting zones more exposed to the China's import competition than others. In the data, these variations are large, as illustrated in the map of Figure 4-3 and quantified in Table 4.3. The key idea is that each US commuting zone specializes in a set of industries but not in all of them (Ellison et al., 2010). Similarly, and centrally to this analysis, China's opening affected a narrow set of industries more heavily and much less some (Autor et al., 2013; Pierce and Schott, 2016a) For example, places specialized in textiles and plastic goods saw sharply larger increases in China's import competition compared to places specialized in the steel, chemical, or paper industries (Autor et al., 2013).

The baseline measure of trade exposure at the CZ level (the instrument) is the local employment-weighted average of changes in the US industry import exposure ratio:

$$\Delta IP_{i\tau}^{CZ} = \sum_j \frac{L_{ijt}}{L_{it}} \times \frac{\Delta M_{j\tau}^{UC}}{M_{j,t_0} - E_{j,t_0} + Y_{j,t_0}}. \tag{4.2}$$

The key component of this measure is $\Delta M_{j\tau}^{UC}$, the change in imports from China in a US manufacturing industry j over the selected period τ (most estimations are performed in stacked annualized decadal differences 1991–1999 and 1999–2011).[11] It is divided by the initial absorption

[8]The methodology draws from research by Autor et al. (2013), and related studies by Autor et al. (2014), Acemoglu et al. (2016), Pierce and Schott (2016a), and Bloom et al. (2016). For a review on the identification strategy and the related literature, see Autor et al. (2016). Dix-Carneiro and Kovak (2017) and Edmonds et al. (2010) use similar strategies based on geographic variations and trade opening in Brazil and India.

[9]UN Comtrade Database 1990–2011.

[10]Autor et al. (2016) provide a comprehensive survey on the factors behind the increase in China's trade.

[11]The year 1991 is the earliest where high-quality disaggregated bilateral trade data are available.

$Y_{j,t_0} + M_{j,t_0} - E_{j,t_0}$ at the baseline year; where M_{j,t_0} is the industry imports, E_{j,t_0} is the industry exports, and Y_{j,t_0} is the industry shipments. The industry-measure tracks export supply shocks from China to US manufacturing output demand in industries where China and the US started to compete after 2001. The industry-level measure is mapped into geographical commuting zones by constructing local industry-employment-weighted sums of industry changes: L_{ijt}/L_{it} is industry j's baseline period share of total employment in CZ i. The variations in the geographical instrument $\Delta IP^{CZ}_{i\tau}$ come from variations in the local industry employment structure in the baseline year.

An alternative measure of trade exposure at the CZ level (the alternative instrument) is analogous but based on China's imports to eight developed countries excluding the US:

$$\Delta IPO^{CZ}_{i\tau} = \sum_j \frac{L_{ijt}}{L_{it}} \times \frac{\Delta M^{OC}_{j\tau}}{M_{j,t_{0-k}} - E_{j,t_{0-k}} + Y_{j,t_{0-k}}}, \tag{4.3}$$

where $\Delta M^{OC}_{j\tau}$ is the change in imports from China in the manufacturing industry j in a set of eight high-income countries that excludes the US.[12] The denominator $M_{j,t_{0-k}} - E_{j,t_{0-k}} + Y_{j,t_{0-k}}$ is defined as above but for the eight other countries.[13] The trade volumes are simply summed over the countries. The employment weights refer to the CZ:s industry employment structure as in the baseline measure. The alternative instrument is motivated by a concern that the baseline US measure can, in part, reflect US-based shocks to US import demand. The alternative instrument aims to capture the supply-component of China's exports to the US, and eventually its impact on US manufacturing industries. The identifying assumption is that the other high-income economies were similarly exposed to China's trade opening and that their industry demand shocks are uncorrelated with each other.[14] Intuitively, the supply component is correlated between the countries, while the demand component is less so. A large literature, surveyed by Autor et al. (2016), highlights that still the main source of variations in China's exports to the US comes from factors internal to China. But the alternative instrument can potentially clean US industry demand shocks from the estimation.

Data on international trade for 1991–2011 come from the UN Comtrade Database. It provides bilateral imports and exports data harmonized at the six-digit HS product level. I match the product-level data to four-digit SIC industries using the crosswalk of Pierce and Schott (2012). The crosswalk assigns 10-digit HS products to four-digit SIC industries (at that level each HS product maps into a single SIC industry). The data from UN Comtrade are at the level of six-digit HS products. At that level some HS products map into multiple four-digit SIC industries. To weight product data to industries, I use US import data at the 10-digit HS level, averaged over 1995–2005.

[12] The countries are Australia, Denmark, Finland, Germany, Japan, New Zealand, Spain, and Switzerland. The specific set of countries is based on data availability: these are the only high-income countries that have bilateral trade data available in 1991 at a level that can be harmonized to HS classification.

[13] The local industry employment data are from 1988 (not 1991) to reduce the error covariance between the dependent and independent variables.

[14] This assumption, made in Autor et al. (2013) and Acemoglu et al. (2016), among others, is rather strong, and unlikely to hold literally. But the alternative instrument can still help overcome some part of the endogeneity issues regarding US import demand shocks.

This process aggregates the four-digit SIC industries to 397 manufacturing industries that all have product codes assigned to them. As in Autor et al. (2013), to match other industry data, I merge a few industries together, resulting in 392 manufacturing industries. All trade amounts are inflated to 2007 US dollars using the Personal Consumption Expenditure (PCE) deflator obtained from the US Federal Reserve. Table 4.3 summarizes the CZ-level changes in exposure to China's imports.

Intuitively, the main estimates come from comparing changes in high-school dropout rates between places with different patterns in manufacturing employment share over time. I focus on the variations in manufacturing employment that come from the exposure to Chinese imports. As argued earlier, these variations plausibly came from outside the system unexpectedly (Autor et al., 2013, 2016). This makes the comparisons between changes in CZ high-school dropout rates and changes in CZ manufacturing employment potentially informative.

To make the comparisons cleaner, I control for a set on baseline characteristics of the places: the baseline manufacturing share of employment, region of the US,[15] employment–to–population ratio, and the population size of the CZ. The baseline manufacturing share control induces comparisons between places that had a similar share of manufacturing employment but saw different declines in it due to differential exposure to China's opening to the world market. This control is important, since variations in the instrument are especially pronounced within the manufacturing sector (see, Tab. 4.4). The regional controls narrow the comparisons to within-region differences, so that the results are not driven by differential trends between regional areas of the US. The controls for employment–to–population ratio and the population size of the CZ narrow further the comparisons to between places with similar employment rates and labor market size. The commuting zone baseline controls are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 periods. In the analysis, treatment is the manufacturing decline, and the comparison group is the otherwise similar places that had a smaller decline in manufacturing. The specifications also include a control for a time-trend.

This research focuses on manufacturing decline, instrumenting it by changes in China's import shares, in contrast with a large research literature initiated by Autor et al. (2013) that studies the labor market consequences of trade with China. That is, the approach of this study creates variations in manufacturing rather than only in trade exposure. From this perspective, Autor et al. (2013) trade exposure estimates depict the reduced form relationship, and my estimates are the IV estimate of interest, scaling up the trade exposure with the induced variations in manufacturing. The estimates are interpretable as the local average treatment effects (LATE) of the manufacturing decline if the China shock works exclusively through its effect on manufacturing employment. Extensive previous research suggest that this is the case (see, Autor et al. 2016, for a review). To be clear, the rapid rise in China's imports to the US had various effects on local labor markets.[16] In the previous literature, these effects have been interpreted to be working through trade exposure's effect on manufacturing industry. But with imperfections in labor and other markets, China's trade shock may have had an independent effect on manufacturing firm revenues, without working

[15]The regional controls indicate nine regional census divisions.

[16]These effects vary from reductions in employment rates to increases crime (Autor et al., 2013; Deiana, 2016).

through changes in the manufacturing employment, translating to incomes and tax revenues that can both affect children's outcomes. While this is unlikely to be qualitatively important, I report both reduced form and IV estimates.

In terms of interpretation, the analysis at the CZ level jointly estimates the reallocation and aggregate demand effects of the manufacturing labor demand shock (as pointed out by Acemoglu et al. 2016). The reallocation effect works through the movement of production factors from the declining sectors to new sectors. The aggregate demand effect multiplies the negative direct and indirect effects of the manufacturing decline stemming from import growth from China. The instrument combines induced employment shifts in both trade-exposed and non-exposed industries. Put simply, the estimates capture the total effect of China-induced manufacturing decline working through many potential channels, including employment-, wage-, and public finance effects, and social and psychological responses within the community.

The IV strategy estimates the local average treatment effect (LATE). It is the effect of treatment on the population of compliers. The compliers are those places that saw a decline in the manufacturing precisely due to China's opening to the world market. The effects of manufacturing decline in these places may differ from the effects in some other places where manufacturing employment declined for some other reason. But this effect in left-behind places hurt by globalization is exactly what this study and many policy makers are interested in (see, for example, Economist 2017). In particular, the LATE may reflect the effect of unexpected manufacturing decline, while the OLS estimates could capture more secular trends. The IV estimates reflect the effect of differential exposure to manufacturing decline, which may differ from the effect from aggregate US manufacturing decline.

Critical threats to the validity of the estimates come from omitted variables correlated with the instruments. A key threat is selective mobility. That is, the empirical strategy essentially considers synthetic cohorts over time in different places. But this idea does not work if the cohorts are significantly unstable over time. Validating this aspect of the identification strategy, evidence from the US suggests that mobility responses to labor market shocks in 1991–2011 have been small and incomplete (Glaeser and Gyourko, 2005). Less educated workers and their families—many of which work in manufacturing and are subject to the largest variations of the treatment—are even less mobile (Notowidigdo, 2011). In particular, investigating mobility and trade shocks, Autor et al. (2013) find little impact of regional trade exposure on changes in mobility. Furthermore, Autor et al. (2014) consider whether workers initially employed in more trade-exposed industries are more likely to change their place of residence, and find little effects.

In summary, the IV strategy constructs plausibly exogenous variations in manufacturing employment between places that without being exposed to the instrument could have had similar trends in educational outcomes. Using this strategy, I can evaluate the effects of trade-induced manufacturing decline on places and children.

4.2.3 Education

A. High-School Dropouts

The main educational outcome is the high-school dropout rate. Data on high-school dropout rates come from the US Census for the years 1990 and 2000, and from the American Community Survey (ACS) for the year 2011.[17] It is defined as the share of civilian 16 to 19 year-old population that is not enrolled in school nor is a high school graduate. The benefit of high-school dropout rate as an outcome is that it captures activity rather than a cumulative stock value. The US Census and ACS report the data at the county level. I match the counts on 16-19 year-old total population and high-school dropout population to the CZ-level using the matching strategy detailed in Dorn (2009), and compute the CZ-level high-school dropout rates. The US Census and ACS are particularly useful data sources for geographical analysis due to their full coverage and large sample size.

For estimation, the main outcome variable is the (annualized) decadal change in the high-school dropout rate in a CZ i over time period τ:

$$\Delta HS_{i\tau} = \Delta \left(\frac{HS_{i\tau}^{16-19}}{POP_{i\tau}^{16-19}} \right) \tag{4.4}$$

where $HS_{i\tau}^{16-19}$ is the number of 16 to 19-year-old residents of the commuting zone (CZ) i that are not in high school nor high-school graduates, and $POP_{i\tau}^{16-19}$ is the population of 16 to 19-year-olds in the same CZ. Most analyses focus on time period τ over 1990–2011.

Table 4.1 shows descriptive statistics for high-school dropout rates in CZs. On average, high-school dropout rate was 10.3 percent in 1991 and decreased to 6.0 percent in 2011. These averages mask large geographical variations in the trends. A map in the Figure 4-3 visualizes the geography of changes in high-dropout rates, and compares it to the changes in the manufacturing employment share. A simple visual comparison suggests that places where high-school dropout rates declined are also places where manufacturing declined.

B. College Mobility

As an alternative measure, I use the college-income gradient developed by Chetty et al. 2014.[18] This outcome variable—college mobility—measures the degree to which a child's college attendance at age 19 is predicted by parental income. It captures one aspect of college access of the young people who were born in a given commuting zone.

College mobility is computed from the restricted access universe of individual tax returns from the U.S. Internal Revenue Service (IRS). In the underlying data, college attendance is defined as an indicator whether the child has a 1098-T form filed on her behalf when she is 18–21. All colleges and

[17] Starting in 2010, the Census stopped using the long form survey and reports education data in the American Community Survey. The American Community Survey measure is computed as a five-year average over 2009–2013. Additional analyses use high-school dropout rates over 1970–1990 from the US Census.

[18] No publicly available US database captures college attendance by the place of birth, previous schooling location, or parental place of residency.

universities, vocational schools, and other postsecondary institutions that are eligible for student aid—are require to file 1098–T forms that report the tuition payments or scholarships received by the student. The 1098-T forms are reported by the universities independently of individual tax returns and plausibly cover the college attendance for all US children. Chetty et al. (2014) document that the tax records capture college attendance quite accurately. The parental income data come similarly from the US tax records, and is defined as the pre-tax adjusted gross income plus tax-exempt interest income and the non-taxable portion of Social Security and Disability (SSDI) benefits. The income measure includes labor earnings, capital income, unemployment insurance, Social Security, and disability benefits, but excludes nontaxable cash transfers, such as food stamps.

This paper uses the public-use summary statistics on intergenerational mobility at the CZ-level provided by Chetty et al. (2014) with an agreement from the IRS. The data are available by CZ for cohorts born between 1984 and 1993.[19] The data include two summary statistics for each CZ and cohort: the estimated slope of a linear equation that predicts college attendance based on parental income, and an intercept. In particular, Chetty et al. (2014) estimate the slope and the intercept of the conditional expectation that a child is attending college given her parents' national income rank for each CZ i and cohort c:

$$C_{jic} = \alpha_{ic} + \beta_c P_{ic} + \varepsilon_{ic} \tag{4.5}$$

where C is an indicator for a child j being enrolled in college at age 19. The slope of the college–income relationship (β_c) measures the degree of relative college mobility in CZ i and for cohort c.[20] The linear conditional expectation fits the data remarkably well (Chetty et al., 2014).

For the analysis of this paper, I construct a measure of "absolute upward mobility" (Chetty et al., 2014) at percentile p in CZ i for cohort c, as the expected probability of attending college for a child who grew up in CZ i with parent who have a national income rank of p: $\bar{c}_{pic} = \alpha_{ic} + \beta_{ic} p_c$. In particular, I focus on the CZ-cohort average of college attendance of children with parents at the 25th percentile in the national distribution, $\bar{c}_{25,ic} = \alpha_{ic} + 0.25\beta_{ic}$.

As the outcome variable, I use an annualized decadal change in the 25th percentile college mobility (CM) between the cohorts born in 1984 and 1993 in the CZ i:

$$\Delta CM_{25,i\tau} = \Delta\left(\alpha_{i\tau} + 0.25\beta_{i\tau}\right). \tag{4.6}$$

The idea behind using changes between cohorts 1984 and 1993 is that most of the manufacturing decline and increase in the China's imports began after 2001, the year the cohort born in 1984 turned 17, and a year before the median starting-age of college. In contrast, the cohort born in 1993 turned 17 in 2010, a year before the end-line of our analysis period. While the control group may also have been affected by the manufacturing decline, the difference between the cohorts captures the

[19] To preserve confidentiality, values for CZ-cohort cells with fewer than 250 observations are omitted.

[20] Note that this reverses the notation of Chetty et al. (2014) to maintain consistency with respect to other notation in this paper.

change in treatment intensity.[21] In line with this idea, I define this measure as the change in college mobility over 1999–2011, the college starting years of each cohort. Tables 4.1 and 4.3, and the map in Figure 4-3 report descriptive statistics on college mobility over 1999–2011.

The drawback of the college mobility data is that it is only available for a single 9-year change. This reduces statistical power and prevents from including controls for time-trends and the pre-period evolution of college mobility. For data confidentiality reasons, the measure is only available for 616 CZs.

4.3 Estimates

The main specification is a stacked first-difference model for annualized decadal changes in the CZ-level variables 1991–2011:

$$\Delta Y_{i\tau}^{CZ} = \alpha_\tau + \beta \Delta MF_{i\tau}^{CZ} + \gamma X_{i0} + e_{i\tau} \tag{4.7}$$

The dependent variable is either ΔHS_{it}, the annual change in the high-school dropout rate in CZ i over time period τ, or $\Delta CM_{25,i}$, the annual change in the 25th percentile college mobility (CM) between the birth cohorts of 1984 and 1993 in CZ i. The term X_{i0} is a set of CZ start-of-period controls; α_τ is the time effect; and $e_{i\tau}$ is the error term. The key explanatory variable in this model is $\Delta MF_{i\tau}^{CZ}$, the annual change in the manufacturing–to–total employment ratio over period τ in CZ i. The coefficient β reveals the impact of manufacturing decline on educational outcomes. The standard errors are clustered by commuting zone to allow for over-time error correlations.

To establish plausibly causal interpretation, I instrument for the decline in manufacturing employment share using the contemporaneous growth of China's imports to the US, $\Delta IP_{i\tau}^{CZ}$, as specified in Section 4.2.2, or alternatively using the growth of China's imports to the eight other high-income countries, $\Delta IPO_{i\tau}^{CZ}$, specified in Section 4.2.2. The variations in the instrument come from variations in local industry employment structure, making some places more exposed to rise in China's exports. Table 4.3 summarizes the CZ-level changes in the key variables: manufacturing share, import exposure, high-school dropout rate, and college mobility.

A. The First Stage

The analysis begins by estimating a first stage relationship between the commuting zone exposure to China's imports and manufacturing decline. The first stage is estimated from stacking changes in CZ manufacturing-to-total employment ratio and exposure to Chinese imports within local industries over the periods 1991–99 and 1999–2011:

$$\Delta MF_{i\tau}^{CZ} = \alpha_\tau + \beta \Delta IP_{i\tau}^{CZ} + \gamma X_{i0} + e_{i\tau} \tag{4.8}$$

[21] Research by Chetty et al. (2014) and Chetty and Hendren (2017a,b) suggests that the effects from the exposure to local conditions come mostly when the children are young. This supports the approach of comparing these cohorts.

The term X_{i0} is a set of CZ-by-sector start-of-period controls, α_τ is the time effect; and $e_{i\tau}$ is the error term. Table 4.4 details the estimates obtained with this approach. The sizable F-statistics for the excluded instruments indicate that regional variations in import exposure have a strong influence on the likelihood of manufacturing decline for CZ:s. The columns 1–3 are estimated without the control for the baseline manufacturing employment share, while the columns 4–6 include that control. Within the CZ:s with the same start-of-period share of manufacturing employment and other baseline controls, the coefficient of trade exposure variable is smaller (-.87 vs. -2.18) but its explanatory power is larger (adjusted R^2 of 0.40 vs. 0.29).

As a visual illustration of the first stage relationship, Figure 4-4 plots the value of the instrument, import exposure as detailed in the Equation 4.2, against the value of the explanatory variable, manufacturing decline as in Equation 4.1, for all US commuting zones over 1991–2011, which is equivalent to the first-stage regression in Table 4.4 but without additional controls and performed in single annual change over 1991–2011. The slope coefficient is –2.80 with standard error 0.21 and t-statistic –13.4. The regression has an R-squared of 0.35, again indicating a relatively strong predictive power of import growth from China for the US manufacturing decline (as also reported by Autor et al. 2013).

B. High-School Dropout Rate Estimates

The OLS and 2SLS estimates of manufacturing decline effects on commuting zone high-school dropout rates 1991–2011 are presented in Table 4.5.

Columns 1–4 present the OLS estimates, progressively including additional baseline controls in the specification. These estimates do not have a causal interpretation, but show a negative relationship between manufacturing decline and high-school dropout rates. In places where manufacturing has declined, high-school dropout rates have declined, too. In columns 1–3, the estimates of the predictive effect vary from –.109 to –.0733 with $p < 0.01$. However, including regional controls for nine US Census divisions make the effect smaller and in the most restrictive model the coefficient is statistically insignificant.

Columns 5–8 present the 2SLS estimates. Column 5 of Panel A considers the relationship between CZ manufacturing decline and changes in CZ high-school dropout rates without additional controls, except for a control for a time trend. The strongly negative and statistically significant point estimate in this column indicates that a 1 percentage point decrease in the manufacturing share of total employment decreases the high-school dropout rate among CZ's 16- to 19-year-old population by .227 percentage points. The OLS and IV estimates are different possibly because the IV estimates capture the effect of unexpected manufacturing decline, while OLS estimates reflect the more predictable secular decline in manufacturing that may have had less impact.

The last three columns of Panel A and Panel B, refine the estimates and explore the robustness, by controlling for the initial manufacturing employment share in a local labor market (Panel B), the initial population (col. 6), the employment-to-population ratio at the baseline (col. 7), and for nine census divisions (col. 8).

By controlling for local manufacturing intensity in Panel B, I allow for differential employment trends in the manufacturing and non-manufacturing sectors. This creates (thought) comparisons between with places with the same manufacturing intensity but saw different changes in it, due to exposure to China's trade. The control for initial population allows for different time trends in local labor markets with different sizes. Similarly, the control for employment-to-population ratio allows for separate trends for labor markets with different levels of activity. The controls for census divisions allow for heterogeneity in regional time trends. The control for the baseline manufacturing employment share has a sizable impact on the estimates. It increases the estimate from –.227 to –.433, without additional covariates. Adding the other covariates has a modest impact on the manufacturing decline coefficient. Among these covariates, the regional controls seem to matter the most. The most restrictive, and preferred, estimate remains sizable and statistically significant at –.366 in column 8 of Panel B.

Taking together the OLS and 2SLS estimates suggests that manufacturing decline is associated with a reduction in high-school dropout rates. In this data, this effect varies between –.16 and –.37 percentage points per a 1 percentage point decrease in the manufacturing share of total employment. In terms of magnitude, the average high-school dropout rate in 1991 was 10.3 percent; and the decline in manufacturing share of total employment across CZs was 7.9 percentage points over 1991–2011. Using the preferred estimate of –.366, this translates to a 2.9 percentage point reduction in the commuting zone high-school dropout rate over 1991–2011—a large but reasonably sized effect.

C. College Mobility Estimates

The OLS and 2SLS estimates of manufacturing decline effects on commuting zone college mobility 1999–2011 are presented in Table 4.6. As described in detail in Section 4.2.3, college mobility is CZ-level average of college attendance of children with parents at the 25th percentile in the national distribution. The college mobility measure comes from Chetty et al. (2014) and is based on the US tax records.

Columns 1–4 present the OLS estimates. They show that, on average, places that saw declines in manufacturing as a share of total employment were also the places that saw increases in college mobility. The predictive effect is smaller but stays statistically and economically significant after controlling for a set of baseline characteristics of these places. The estimates of the predictive effect vary from .397 to .254 with $p < 0.05$.

Columns 5–8 present the 2SLS estimates. In Panel A, without additional controls, the estimate in the column 5 considers the relationship between CZ manufacturing decline and changes in CZ college mobility. Consistent with the results of the high-school dropout rate analysis, the 2SLS estimate in column 5 implies a positive and statistically significant effect from manufacturing decline to college mobility. In particular, the estimate in this column indicates that a 1 percentage point decrease in the CZ's manufacturing share of total employment increases the CZ average of college attendance of children with parents at the 25th percentile in the national distribution by .36 percentage points. Controlling for the baseline population size and employment-to-population ratios leaves the 2SLS estimates largely unchanged. Including the nine regional census indicators makes the estimate insignificant, but keeps its sign unchanged and the magnitude in the ballpark.

In Panel B, focusing on the variations within a set of places with similar manufacturing start-of-period share of total employment, the estimates are not anymore statistically significant. However, most coefficients do have the same sign and, while smaller, fit into the range of the estimates of Panel A. The college mobility variable covers only a one observation per CZ. A plausible interpretation is that including the manufacturing share control leaves too few degrees of freedom to produce precise estimates.

Although, in general, the estimates highlight a negative descriptive relationship between manufacturing share of employment—working class jobs—and college mobility, the most restrictive causal estimates are inconclusive.

In terms of interpretation, a drawback of the college mobility measure is that it leaves a few possibilities for the mechanism driving the increases (or decreases) in it.[22] The simple case is that relatively poorer children enroll more in college. However, the focus on the national distribution creates a complication when looking at changes over time. A decline in local income moves the residents left in the national income distribution. But if income is not an important determinant of college access in that place, this decline in incomes translates to an increase in the college mobility measure: now poorer children (that were previously rich) are more likely to go college. The drawback aside, supporting a non-mechanical interpretation, Chetty and Hendren (2017a,b) provide evidence that the given college mobility rates of a CZ are largely interpretable as causal effects of the place. While the most restrictive estimates are inconclusive, the research of this paper suggest a potential causal chain from lost manufacturing jobs to a place that provides higher college access to poor children. To establish or dispute the causal chain, more research is needed.

4.4 Robustness

A. Pretrends

US high-school dropout rate has been declining since the 1970s,[23] and manufacturing as a share of employment has also trended downward since the 1950s.[24]A visual inspection of the maps in Figure 4-3 suggest that in the period 1991–2011, these trends tended to be stronger in the same places. This association could, however, be a result of a long-standing secular trend. The correlation this study documents between declining manufacturing share of employment and contemporaneous declines in high-school dropout rates during 1991–2011 could potentially predate the recent decline in manufacturing. In that case, the estimates would likely overstate the impact of manufacturing decline in the current period. To address this concern, I include measures of pretrends in high-school

[22]This complication is not highlighted in Chetty et al. (2014). Unfortunately, the raw data on college attendance by place of birth, or similar data, are not available.
[23]Source: US Census.
[24]Source: Community Business Patterns.

dropout rates in Table 4.7, specifically two terms for the change in the CZ high-school dropout rates, measured over the intervals 1970–80 and 1980–1990.[25]

Formally, the pre-trend controls mean including lagged dependent variables to the stacked first-difference specification:

$$\Delta Y_{i\tau}^{CZ} = \alpha_\tau + \beta \Delta MF_{i\tau}^{CZ} + \gamma X_{i0} + \delta_1 \Delta Y_{i1980\text{–}90}^{CZ} + \delta_2 \Delta Y_{i1970\text{–}80}^{CZ} + e_{i\tau} \tag{4.9}$$

Table 4.7 replicates the main set of results on high-school dropout rates but including the pretrends. The pretrend variables have no important effect on the magnitude or precision of the coefficient of interest: the estimates are close to that found in the main Table 4.5. The measured effects are slightly larger, increasing from −.366 to −.418 with $p < 0.01$ in the preferred and most restrictive specification. However, this hints that even with the IV strategy, there is some temporal dependence left in the local high-school dropout series.

B. Falsification Test

As a falsification test, Table 4.8 reports results from a 2SLS regression of changes in high-school dropout rates in earlier decades on the instrumented manufacturing decline between 1999 and 2011. For the identification strategy, it would be a concern if *future* declines in CZ manufacturing due to China's trade opening predicted *past* changes in local high-school dropout rates—in the time periods before China had affected US manufacturing. Operationally, I estimate a set of models:

$$\Delta MF_{i\tau}^{CZ} = \alpha_i + \beta \Delta IP_{i1999\text{–}2011}^{CZ} + \gamma X_{i0} + e_{i\tau}, \tag{4.10}$$

where τ takes four different values: 1970–80, 1980–90, 1990–2000, and 1999–2011.

In Panel A, the first row performs the estimation without additional controls. The rows 2–4 go through combinations of regional controls and the controls for baseline share of manufacturing in total employment. The results from the specifications that include baseline controls, either the regional controls or manufacturing share, show largely that future instrumented manufacturing decline does not predict past changes in high-school dropout rates. Conversely, the estimate is large and significant in the contemporaneous period 1999–2011 where it should be. Adding demographic covariates keep the estimates essentially unchanged (not reported). This pattern of results is consistent with the identifying assumption that the within-industry and CZ correlation between declining manufacturing employment and import penetration from China in 1991–2011 that seems to translate to reductions in high-school dropout rates, originates from trade shocks rather than long-term secular trends between manufacturing employment and high-school dropout rates.

However, the specifications that do not include any covariates show some evidence of temporal dependence in the high-school dropout rate series. The predictive effect is visible for the 1970–1980

[25] Data on college mobility are only available over 1999–2011 and thus does not allow for testing pretrends in that outcome variable.

period.[26] But including heterogeneous regional trends makes this effect disappear. This pattern of findings suggests that the regional baseline covariates are necessary for the identification of the empirical results. In either pretrend analysis or falsification test, the most restrictive manufacturing employment control does not make a difference.

C. Alternative IV

A reasonable concern is that the measured US imports from China, used to construct the main instrument, could be correlated with domestic US demand shocks rather than reflecting external supply factors external to the US labor and product markets, possibly resulting in biased estimates. As detailed in Section 4.2.2, an alternative instrument uses China's import growth in eight other high-income countries as detailed in Equation 4.3. The idea is that the other high-income face a similar supply shock from China, while are possibly subject to different idiosyncratic industry-specific demand shocks. Intuitively, China's trade flows to other countries than the US are plausibly less determined by factors internal to the US and more by factors related to China's opening to the world market.

The alternative instrument is explored by estimating the main specification,

$$\Delta Y_{i\tau}^{CZ} = \alpha_\tau + \beta \Delta MF_{i\tau}^{CZ} + \gamma X_{i0} + e_{i\tau} \tag{4.11}$$

but instrumenting the changes the manufacturing as a share of total employment, $\Delta MF_{i\tau}^{CZ}$, with the contemporaneous change in China's imports elsewhere, $\Delta IPO_{i\tau}^{CZ}$.

Table D.1 reports the results from the alternative IV estimation both for high-school dropout rates 1991–2011 and college mobility 1999–2011. The point estimates are almost identical to the main estimates of Tables 4.5 and 4.6. The first stage relationship is equally strong with F-statistic 155.0 and adjusted R^2 of .22 without baseline controls, and F-statistic 71.7 and adjusted R^2 of .39 with a full set of baseline controls.

D. Reduced Form Estimates

Interpreting the IV estimates as the effect of manufacturing decline requires assuming that China's trade exposure affected educational outcomes exclusively through its effect on manufacturing employment. Recall, that the variations in the instrument come exclusively from differential manufacturing industry compositions between places. Therefore, from this perspective, the assumption is not unreasonable (see Section 4.2.2 for further discussion). But although most of the effect are likely to translate through the manufacturing industry, the increased competition could affect the community and thus education through manufacturing industry profits and reduced demand for suppliers rather than employment. Based on other studies these income- and public finance effects would likely reduce human capital investment, and bias the estimates downwards (see, for

[26] A significant predictive effect for 1970–1980 from China's imports to manufacturing employment is similarly found in Autor et al. (2013), p. 2135.

example, Davis and von Wachter 2011, for a review).[27] For conceptual clarity, in addition to main IV estimates, I provide the reduced form estimates from trade to educational outcomes, both for high-school dropout rates and college mobility.[28] In particular, I estimate the following model:

$$\Delta Y_{i\tau}^{CZ} = \alpha_\tau + \beta \Delta IP_{i\tau}^{CZ} + \gamma X_{i0} + e_{i\tau} \tag{4.12}$$

where the dependent variable is either ΔHS_{it}, the annual change in the high-school dropout rate in CZ i over time period τ, or $\Delta CM_{25,i}$, the annual change in the 25th percentile college mobility (CM) between the cohorts of 1984 and 1993 in CZ i. The term X_{i0} is a set of CZ-by-sector start-of-period controls; α_τ is the time effect; and $e_{i\tau}$ is the error term. The explanatory variable in this model is $\Delta IP_{i\tau}^{CZ}$, the annual change in exposure to Chinese imports within local industries over period τ in CZ i. The coefficient β reveals the impact of trade exposure on educational outcomes. The standard errors are clustered by commuting zone to allow for over-time error correlations. Table D.2 presents the results.

In Columns 1 and 2, the reduced form estimates have the same signs and similar magnitudes than the IV estimates for both high-school dropout rates and college mobility. As before, including baseline control for the manufacturing employment share makes the college mobility coefficient insignificant. However, the coefficient is still large, positive, and significant when including all other baseline controls.

In Columns 3 and 4, the import exposure instrument itself is instrumented with the alternative instrument constructed from Chinese imports to eight other high-income countries, as in Autor et al. (2013). This specification produces larger results, the estimates increase almost by a factor of two. This suggests that quantitative results are somewhat sensitive to the choice of particular instrument, but qualitatively show the same pattern.

E. Log-Log Specification and Baseline Education

So far, the analysis has adjusted for the baseline differences by considering first differences of the variables, controlling for some baseline characteristics of the places, and using the IV strategy. A concern might be still that the places with initially higher high-school dropout rates might have larger response to manufacturing decline. And these places could be the same places where manufacturing declined due to exposure to China's trade. This could bias the results upwards. Now, I consider two extensions to address this. The following discussion focuses on the high-school dropout rates, because the college mobility results with this research design were inconclusive.

First, I estimate the main specification in logarithms:

$$log(\Delta HS_{i\tau}^{CZ}) = \alpha_\tau + \beta log(\Delta MF_{i\tau}^{CZ}) + \gamma X_{i0} + e_{i\tau}, \tag{4.13}$$

[27] The reduced form estimates are also less sensitive to measurement error.

[28] Note that simply controlling for, say changes in the unemployment rate would be bad control, since those changes would most likely be caused by the manufacturing industry exposure to China's trade.

with the same notation, variables, and instrumentation as earlier. This specification considers relative changes in manufacturing and the high-school dropout rate. Table D.3 reports the results. The estimates are similar in sign, significance, and magnitude to the main results that were estimated in percentage points. The estimate −.498 in column 2 means that a 1 percentage (relative) decline in the manufacturing share of total employment decreases the high-school dropout rate among a CZ's 16- to 19-year-old population by .498 percentages.

Second, I control for the baseline high-school dropout rate:

$$\Delta HS_{i\tau}^{CZ} = \alpha_\tau + \beta \Delta MF_{i\tau}^{CZ} + \gamma X_{i0} + HS_{it}^{CZ} + e_{i\tau}, \tag{4.14}$$

where HS_{it}^{CZ} controls are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 period. Now initial high-school dropout rate is used both to compute the dependent variable and as a control variable. While this is a different model than the main specification, the key idea is making treated and control units comparable on lagged outcomes.[29] Table D.3 reports results from this specification. The estimates show more dispersion, but are in line with the earlier results. The preferred estimate with full set of baseline controls is almost unchanged.

4.5 Exploring the Mechanism

A. Rural vs. Urban

Are the effects of manufacturing decline on children's education similar around the US, or are the effects different in rural versus urban America? To study this, I estimate interactions between CZ manufacturing decline and the CZ being located in a rural part of the US. The US Census measures the share of rural population in each US county based on where people work (Ratcliffe et al., 2016). I match this data to CZs weighting by the population of each county. I define a rural CZ as a place where more than 50 percent of the population lives in a rural setting, and compute an indicator variable for it. I explore different thresholds up until 90 percent, with no large impact on results.

Table 4.9 presents the results for rural-urban interaction analysis. The rural-interaction coefficient is small, insignificant, and has only a minor effect on the main coefficient. While not decisive, the results suggest the rural-urban distinction does not play a key rule in channeling the effect. This is a substantial finding: many commentators feel that manufacturing decline and the issues associated with it are particularly an issue of rural America (Economist, 2017). But the implications for children's education appear similar in both.

B. Correlates of the Intergenerational Effects

What characterizes the places where manufacturing decline tends to lead to lower high-school dropout rates? One would expect that some place-based characteristics would matter. For example, places with high income- and racial segregation might not be able to channel students to high-school

[29] Imbens and Wooldridge (2009) provide an informed discussion on the details of this.

after a decline in manufacturing employment. Again, places equipped with generous educational resources might see larger decreases in high-school dropout rates after manufacturing jobs have declined. But the following results show that neither is true—the opposite is.

To produce these results, I estimate interactions between CZ manufacturing decline and a large set of CZ's baseline community factors that have been discussed in the sociology and economics literature, such as segregation and inequality. Because most of these factors are relatively stable over time, and I only have data for essentially one period, I focus on cross-sectional characteristics. The estimation is in stacked annualized decadal differences over 1991–2011:

$$\Delta HS_{i\tau}^{CZ} = \alpha_\tau + \beta_1 \Delta MF_{i\tau}^{CZ} + \beta_2(\Delta MF_{i\tau}^{CZ} \times K_i) + \beta_3 K_i + \gamma X_{i0} + e_{i\tau} \tag{4.15}$$

The outcome variable is ΔHS_{it}, the annual change in the high-school dropout rate in CZ i; K_i is the time-invaryING interacted community variable included in each model one at a time, X_{i0} is a set of CZ start-of-period controls; α_τ is the time effect; and $e_{i\tau}$ is the error term. Again, the main explanatory variable in this model is $\Delta MF_{i\tau}^{CZ}$, the annual change in the manufacturing–to–total employment ratio over period τ in CZ i. The key parameter of interest in this model is β_2, the coefficient of the interaction term. With two endogenous variables, I instrument for the decline in manufacturing employment share using the contemporaneous growth of China's imports to the US, $\Delta IP_{i\tau}^{CZ}$, as specified in Section 4.2.2, and with the interaction term between the fixed community variable and China's imports, $\Delta IP_{i\tau}^{CZ} \times K_i$. The analysis of community characteristics is limited to the 2SLS estimates on high-school dropouts since the estimates for college mobility are considerably more sensitive to specific controls and regional trends.

Tables D.4 and D.5 describes the set of interacted variables and their sources. The data on local factors was compiled by Chetty et al. (2014). The authors provide a comprehensive overview on variable definitions and measurement.

Two main results emerge (Tab. 4.10). First, segregation and share of black population strongly interact with the positive effects of manufacturing decline on education. That is, the effects of manufacturing decline are largest in the areas with high segregation and in those with larger African American populations. This is true for several different measures of segregation. While the risk of false rejections of the null is present with multiple testing, the fact that many different measures of segregation produce a similar result supports this finding.

Second, educational resources—student-teacher ratio and school expenditure per student—do not significantly interact with the main effects. An exception is the number of colleges in CZ, which has a predictive effect of making the manufacturing effect smaller. Additionally, some other insignificant coefficients are noteworthy. For example, religion and social capital (Putnam 1995; measured as activities related to civil society) are both strong predictors of upward income mobility (Chetty et al., 2014). However, they do not interact with the effects of manufacturing decline on children's education.

What the main results seem to suggest is that manufacturing jobs—or a broader working-class community—keep a pathway open for teenagers to drop out of high school. When those jobs decline, the pathway declines, too. Now this effect appears to be stronger in more segregated places. Perhaps the factors behind segregation, or segregation itself, support the pathway. Sociological work by Willis (1977), among others, supports this hypothesis.

Compared to the previous literature on the determinants of children's outcomes, these interactions show significance for very different factors than, for example, for which Chetty et al. (2014) find positive predictive power. In particular, segregation strongly correlates with low upward mobility in a cross section of CZs. And educational resources strongly predict high mobility (Chetty et al., 2014). Perhaps surprisingly, manufacturing decline has the highest effect in those places that on average fail to produce upward mobility of income.[30]

The results suggest that trade-induced manufacturing decline leads to lower high-school dropout rates—especially so in segregated places and those with larger share of black residents. As matter of correlation, local investment level in schooling does not predict the effect. These are new and puzzling findings.

4.6 Discussion

What explains the main results? The standard Becker (1964) human capital investment model compares the marginal costs and benefits of education. The key primitive is the economic returns to education. The model has lead many labor economists to argue that educational investment would increase vis-à-vis increasing income inequality (Ananat et al., 2017). The simplest model argues that manufacturing workers' children would notice that manufacturing no longer offers stable jobs, and would obtain higher education than what their parents had.[31]

The sociological account of Willis (1977) offers a complementary view, highlighting how children inherit occupations and class from their parents. It argues that predominantly working class communities—places where the share of manufacturing employment is high—help children embrace an anti-school mentality and prepare them for low-education working-class employment. For example, Willis (1977) argues, working-class fathers may act as role-models to their children and through that channel affect the children's educational choices. Extrapolating from Willis (1977)'s observations in working-class communities, a decline in manufacturing could lead to a decrease in such role models and translate to an increase in children's education. The results suggests that this may have been the case.[32] Willis' theory could also help reconcile the interaction effects between

[30]In the cross-sectional evidence of Chetty et al. (2014), the share of manufacturing employment only weakly and negatively correlates with upward mobility.

[31]Goldin and Katz (2000) describe the lack of manufacturing jobs and its consequences on education—in the US prairie states of 1910: "Youths in these states could not have worked in industry, for there was scant manufacturing – –. And although many farmers would have preferred that their children remain on the land, most knew it would prove impossible for all but one. The best they could do was to endow their children with education to be mobile."

[32]Commenting on the work in the line of Willis (1977), Vance (2016), pp. 246, confirms this observation in his personal memoir of growing up in rural America: "working-class boys like me do much worse in school because they view schoolwork as a feminine endeavor." Vance (2016), pp. 244, also suggests—from his own experience and observations—that psychological and social factors could be much more important than traditional economic factors: "My elementary and middle schools were entirely adequate – –. I had Pell Grants and government subsidized low-interest student loans that made college affordable. The real problem for so many of these kids is what happens (or

local segregation and manufacturing decline. Perhaps working-class culture was stronger in the areas with higher segregation.

In contrast, available evidence on job loss and income shocks indicates that negative shocks lead to negative effects on children's education (Davis and von Wachter, 2011). This is clearly seen in the scarring effects of parental job loss observed by Oreopoulos et al. (2008). That literature suggests that education investment benefits from the resources that are available to the child. In addition to parental effects, Wilson (1996) in "*When Work Disappears*" and earlier Whyte (1943) point to the loss of jobs, fuelled by decline in manufacturing, as a driver of social anomie and community-level distress in poor neighborhoods and increasing childhood poverty (Autor et al., 2017).

These two lines of thinking—positive factors of Becker (1964) and Vance (2016), and the negative factors of Oreopoulos et al. (2008) and Wilson (1996)—highlight a tension between *income* and *opportunity* shocks. In the empirical literature, this tension is perhaps clearest in the case of Indian casino openings. Federal legislation in 1988 allowed Indian tribes to open casinos in many states, leading to the opening of nearly new 400 casinos in the US. It had both components: the income shock—the casino openings initiated a government transfer scheme giving a portion of the casino profits to individuals with preexisting American Indian status—and the opportunity shock—change in the local employment opportunities. The two components appear to have had opposite effects (Akee et al., 2010; Evans and Kim, 2008). Akee et al. (2010) find that children in the households affected by the the government income transfer program had higher levels of education in their young adulthood. In contrast, Evans and Kim (2008) find that—within the same communities—young adults responded to the increased employment and wages of low-skilled workers by dropping out of high school and reducing college enrollment rates. This was despite presence of the income transfer scheme and additional college tuition subsidy programs of many tribes. This paper and a larger literature suggest that both results could be true at the same time (see, for example, Black et al. 2005; Atkin 2016; Shah and Steinberg 2017).[33]

In summary, while places experiencing manufacturing job losses face reductions in the monetary and social resources available to children, perhaps changing incentives and social structure counteract that effect to produce the positive results on education found in this paper. This adjustment, however, may be incomplete. This paper contributes to the growing evidence on the interplay between local labor-market conditions and educational decisions (Atkin, 2016; Shah and Steinberg, 2017), going beyond the direct effects of parental job loss.

Looking from a different perspective, a large literature discusses regional divergence of the US (Ganong and Shoag, 2017), and why we tend to observe a permanent decline in a place hit by a negative economic shock (Blanchard and Katz, 1992; Dix-Carneiro and Kovak, 2017). Previous scholars have investigated the role of social distress (Wilson, 1996; Ananat et al., 2017), imperfect

doesn't happen) at home."

[33]The tension between income and opportunity shocks has not been generally clear in the research literature. For example, Akee et al. (2010) suggest that the contrary results by Evans and Kim (2008) arise from identification issues from focusing on community-level variations (this is possible). But local employment opportunity effects exist at the community level and therefore the natural observation unit is the community, not the individual.

mobility, declining housing prices, generous social welfare payments (Ganong and Shoag, 2017), and human capital externalities (Dix-Carneiro and Kovak, 2017). The evidence is inconclusive. This paper explores a new channel: human capital investment of the next generation. Local job destruction could lead the youth off the path to high school and college. In the long run, this could lead to lower local productivity and long-term decline. But the results of this research suggest otherwise: a decline in formal education after an economic shock does not seem to be a channel for local long-term decline and regional divergence. Something else is.

4.7 Conclusion

This paper provides new evidence on the impact of manufacturing decline on children. To do so, it considers variations in local employment structure—characterizing left-behind places and lost manufacturing jobs—high-school dropout rates, and college access in the US over 1990–2010. To establish causal inference, the paper uses variations in trade exposure from China following its entry to the WTO as an instrument for local manufacturing declines in the US.

The results suggest that negative shocks to manufacturing labor demand, measured at the local labor market level, had large positive effects on children's education, decreasing high-school dropout rates and possibly increasing college access. The magnitudes of the estimates suggest that for every 3-percentage-point decline in manufacturing as a share of total employment, high-school dropout rate declined by 1 percentage point. These findings contrast with the literature on job loss that has emphasized negative effects from economic shocks on children. The results are consistent with the idea that the manufacturing decline increased returns and decreased opportunity costs of education, and with sociological accounts linking working-class environment and children's education. These effects perhaps counteract the negative effects from income loss. The effects are largest in the areas with high segregation and in those with larger African American populations. This set of findings is new—and a first step in quantifying the intergenerational effects of lost manufacturing jobs due to technological change and globalization.

Children face the collateral damage from the adults' world. And the long-run consequences depend on them. That's why this research matters.

Main Figures and Tables

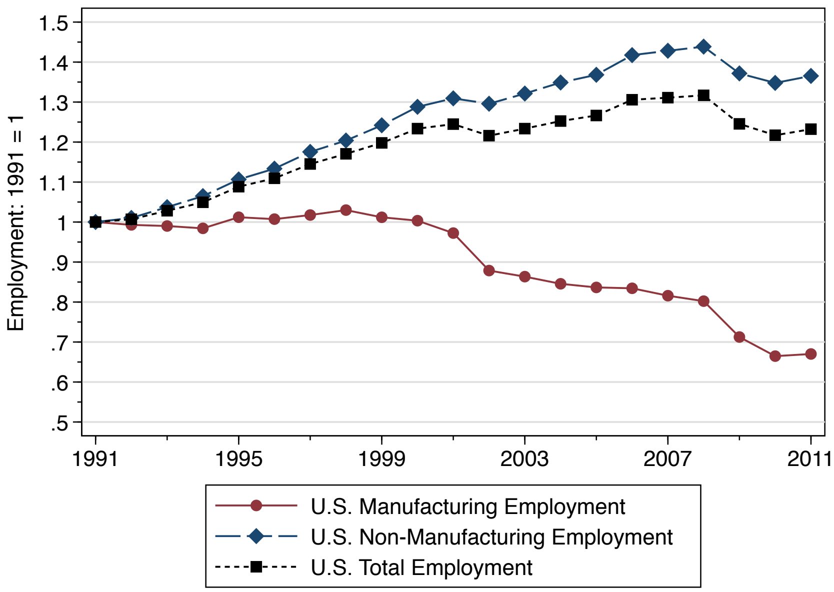

Figure 4-1: Changes in US manufacturing and non-manufacturing employment, 1991–2011. Employment data are normalized to 1991. Source: CBP.

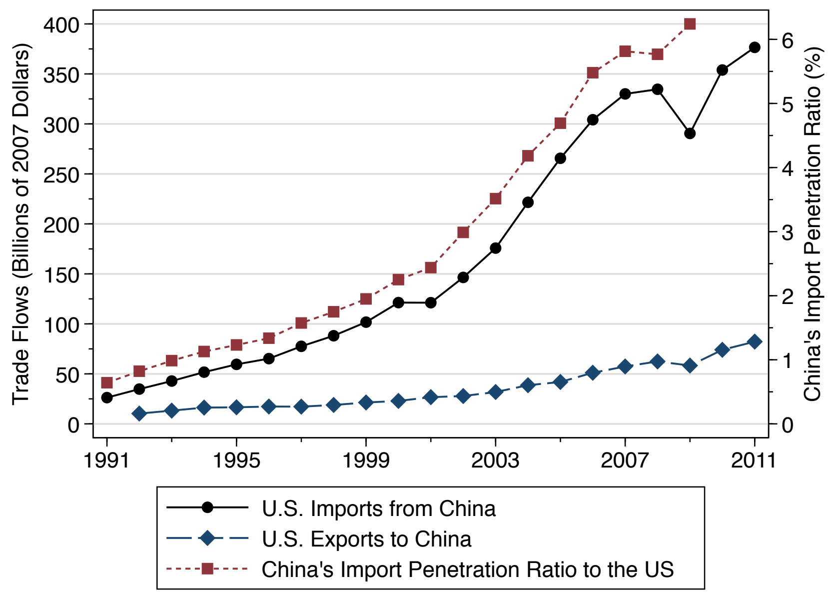

Figure 4-2: US-China bilateral trade flows, 1991–2011. Source: UN Comtrade Database. Trade volumes are deflated to 2007 US dollars using the PCE price index. China's import penetration is defined as China's manufacturing imports to the US divided by US domestic manufacturing output plus imports minus exports. Export data are available only from 1992 onward. The import penetration ratio series ends in 2009 because the NBER-CES Manufacturing Industry Database ends in 2009.

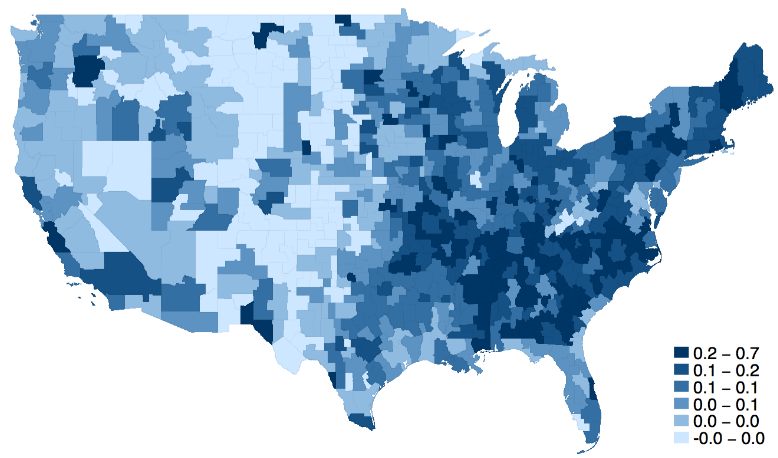

(a) Decline in Manufacturing Share of Employment (–) (CBP, US Census).

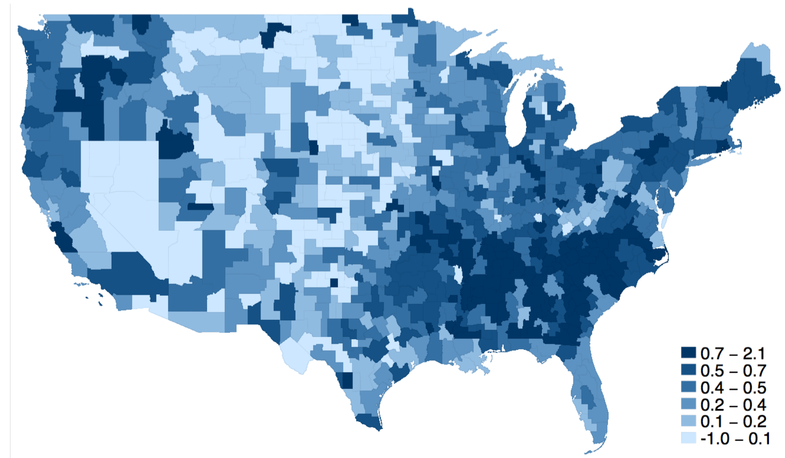

(b) Increase in Import Exposure from China (+) (UN Comtrade).

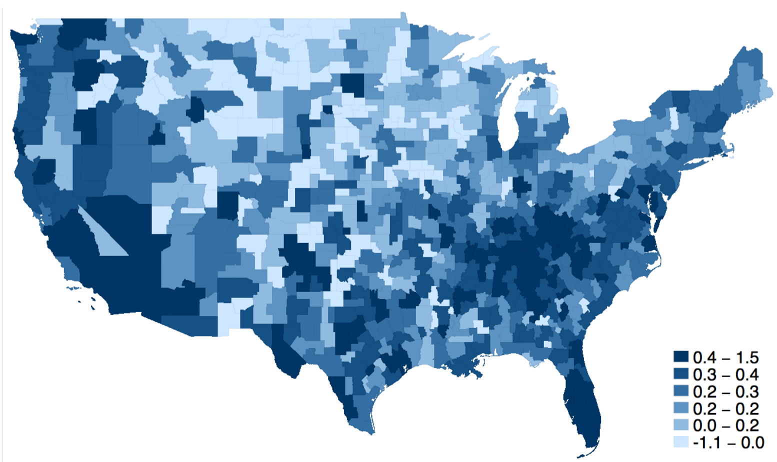

(c) Decline in High-School Dropout Rate (–) (US Census, ACS).

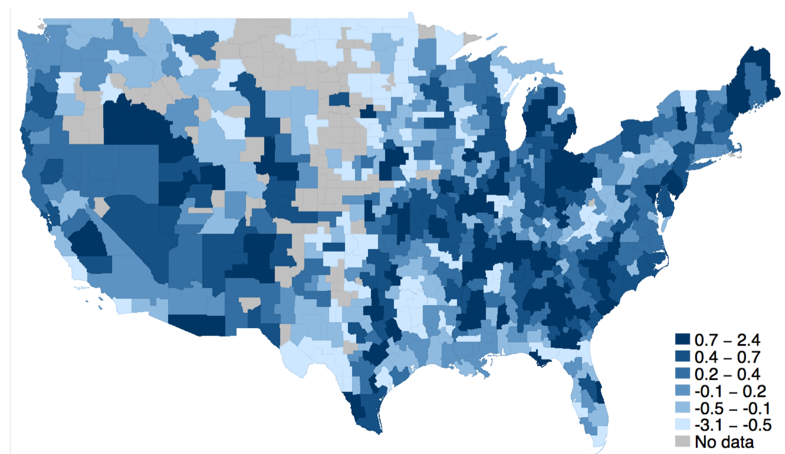

(d) Increase in College Mobility (+) (US Tax Records via Chetty et al. 2014).

Figure 4-3: Maps. All variables are in 100 × annual changes 1991–2011, except college mobility 1999–2011. The (+) and (–) signs indicate whether the heat map refers to increases or decreases in the variable. The variables are constructed as detailed in text.

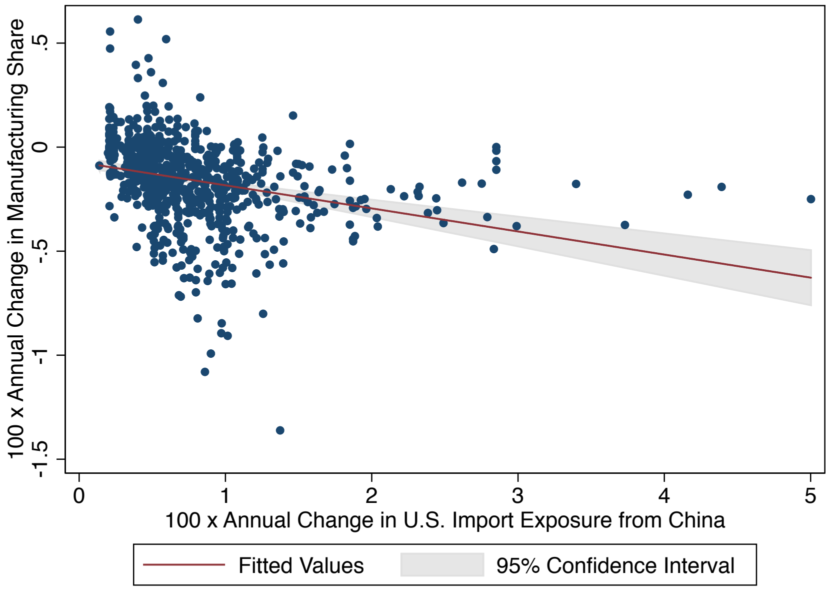

Figure 4-4: First-stage regression, 1991–2011. Each point represents a commuting zone (N = 722). Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. The annual change in commuting zone exposure to Chinese imports is a weighted average of changes in US import exposure in 392 four-digit manufacturing industries, where the weights are start-of-period employment shares within the commuting zone. Imports are deflated to constant dollars using the PCE price index. Lines are fitted by OLS regression. The 95 percent confidence interval is based on standard errors clustered on 722 commuting zones. The slope coefficient is –2.80 with standard error 0.21 and t-statistic –13.4; the regression has an R-squared of 0.35.

Table 4.1: Descriptive Statistics for Manufacturing, Employment, and Population in CZs.

		Mean	Std. Dev.	Min	Max
Manufacturing–to-Total Employment Ratio (%)	1991	21.9	12.8	.10	61.4
	1999	19.3	11.3	.13	57.7
	2011	13.9	8.7	.26	51.3
Employment-to-Population Ratio (%)	1991	42.1	10.6	11.0	76.8
(Working age)	1999	48.01	11.8	16.3	83.0
	2011	44.9	10.4	16.5	80.0
Population	1991	350,000	.95 M	1311	10.4 M
(Total)	1999	380,000	1.04 M	1213	16.6 M
	2011	430,000	1.16 M	1017	18.1 M

Notes: N = 722 commuting zones. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. Working-age population is those between the ages of 15 and 64.

Table 4.2: Descriptive Statistics for Education in CZs.

		Mean	Std. Dev.	Min	Max
High-School Dropout Rate	1990	10.3	4.2	.36	31.7
	2000	9.2	3.9	.41	22.4
	2011	6.0	3.3	.38	30.2
College Mobility	2002	32.5	8.4	13.7	61.2
	2011	33.2	7.7	12.1	58.2

Notes: N = 722 commuting zones for high-school dropout rates, 616 for college mobility, The variables are expressed in percentages. High-school dropout rate is computed from the US Census for 1990 and 2000, and from the ACS for 2011 as a five-year average. College mobility is CZ-level average of college attendance of children with parents at the 25th percentile in the national distribution of income. The college mobility measure comes from Chetty et al. (2014) and is based on the US tax records. The years 2002 and 2011 refer to the standard college-starting years of cohorts born in 1984 and 1993.

Table 4.3: Changes in Commuting Zone Manufacturing Share, Import Exposure, High-School Dropout Rate, and College Mobility.

	1991–99				1999–2011			
	Mean	Std. Dev.	Min	Max	Mean	Std. Dev.	Min	Max
Δ in Manufacturing-to-Total Employment Ratio	-.31	.57	-5.34	3.14	-.45	.43	-2.83	.97
Δ in Exposure to China's Imports	.06	.08	.00	.95	.09	.09	.00	.69
Δ in High-School Dropout Rate	-.10	.29	-1.96	1.43	-.30	.35	.35	1.58
Δ in College Mobility	–	–	–	–	.11	.75	-3.12	2.38

Notes: N = 1444 = 2 × 722 commuting zones. All variables are 100× annual change in the measure. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. The annual change in commuting zone exposure to Chinese imports is a weighted average of changes in US import exposure in 392 four-digit manufacturing industries, where the weights are start-of-period employment shares within the commuting zone. Imports are deflated to constant dollars using the PCE price index. High-school dropout rate is computed from the US Census for 1990 and 2000, and from the ACS for 2011 as a five-year average. College mobility is CZ-level average of college attendance of children with parents at the 25th percentile in the national distribution of income. The college mobility measure comes from Chetty et al. (2014) and is based on the US tax records. The annual change in college mobility in 1999–2011 refers to the annual change in college mobility between cohorts born in 1984 and 1993.

Table 4.4: The First Stage: Estimates of China's Import Effects on Commuting Zone Manufacturing Decline over 1991–2011.

Manufacturing Share	(1)	(2)	(3)	(4)	(5)	(6)
Commuting zone import exposure	-2.58***	-2.54***	-2.18***	-1.25***	-.91***	-.87***
	(.23)	(.24)	(.24)	(.24)	(.23)	(.22)
F-Statistics	117.4	76.5	36.1	216.4	180.6	70.5
Adjusted R^2	0.21	0.21	0.29	0.32	0.33	0.40
Time effect controls	–	Yes	Yes	–	Yes	Yes
Baseline controls	–	–	Yes	–	–	Yes
Manufacturing share baseline	–	–	–	Yes	Yes	Yes

Notes: First stage regression. Each column reports results from stacking changes in commuting zone manufacturing-to-total employment ratios and in exposure to Chinese imports within local industries over the periods 1991–99 and 1999–2011. The dependent variable is the annual change in the manufacturing-to-total employment ratio (N = 1,444 = 722 commuting zones × 2 periods). The explanatory variable an employment-weighted average of annualized changes in exposure to Chinese imports within local industries, as detailed in the text. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. Baseline controls include population counts, employment-to-population ratios, and region controls for nine regional census divisions. The commuting zone baseline controls, including the manufacturing share control, are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 periods. Standard errors are clustered by commuting zone.

$*p < 0.10$

$**p < 0.05$
$***p < 0.01$

Table 4.5: OLS and 2SLS Estimates of Manufacturing Decline Effects on Commuting Zone High-School Dropout Rates 1991–2011.

High-School Dropout Rate	OLS				2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Excluding Manufacturing Share Control at the Baseline								
Commuting zone manufacturing decline	-.109***	-.107***	-.107***	-.044**	-.227***	-.228***	-.232***	-.162***
	(.018)	(.018)	(.018)	(.021)	(.031)	(.030)	(.030)	(.034)
Time effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Population counts at the baseline	–	Yes	Yes	Yes	–	Yes	Yes	Yes
Employment-to-population ratios at the baseline	–	–	Yes	Yes	–	–	Yes	Yes
Census division indicators	–	–	–	Yes	–	–	–	Yes
B. Including Manufacturing Share Control at the Baseline								
Commuting zone manufacturing decline	-.082***	-.075***	-.0733***	-.030	-.433***	-.416***	-.415***	-.366**
	(.021)	(.021)	(.021)	(.025)	(.131)	(.132)	(.130)	(.125)
Time effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Population counts at the baseline	–	Yes	Yes	Yes	–	Yes	Yes	Yes
Employment-to-population ratios at the baseline	–	–	Yes	Yes	–	–	Yes	Yes
Census division indicators	–	–	–	Yes	–	–	–	Yes

Notes: Each column reports results from stacking changes in commuting zone high-school dropout rates and declines in manufacturing-to-total employment ratios over the periods 1991–99 and 1999–2011. The dependent variable is the annual change in the high-school dropout rate (N = 1,444 = 722 commuting zones × 2 periods). The manufacturing decline is instrumented with the commuting zone import exposure from China's imports. The instrument is an employment-weighted average of annualized changes in exposure to Chinese imports within local industries, as detailed in the text. High-school dropout rate is computed from the US Census for 1990 and 2000, and from the ACS for 2011 as a five-year average. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. The commuting zone baseline controls are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 periods. Census division indicators control for nine regional census divisions. Standard errors are clustered by commuting zone.

$*p < 0.10$

$**p < 0.05$

$***p < 0.01$

Table 4.6: OLS and 2SLS Estimates of Manufacturing Decline Effects on Commuting Zone College Mobility 1999–2011.

College Mobility	OLS				2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Excluding Manufacturing Share Control at the Baseline								
Commuting zone manufacturing decline	.397***	.409***	.407***	.316***	.360***	.378***	.329***	.232
	(.076)	(.076)	(.076)	(.082)	(.124)	(.123)	(.124)	(.142)
Population counts at the baseline	–	Yes	Yes	Yes	–	Yes	Yes	Yes
Employment-to-population ratios at the baseline	–	–	Yes	Yes	–	–	Yes	Yes
Census division indicators	–	–	–	Yes	–	–	–	Yes
B. Including Manufacturing Share Control at the Baseline								
Commuting zone manufacturing decline	.387***	.361***	.395***	.254**	.236	.138	.122	-.122
	(.101)	(.101)	(.102)	(.103)	(.278)	(.285)	(.283)	(.301)
Population counts at the baseline	–	Yes	Yes	Yes	–	Yes	Yes	Yes
Employment-to-population ratios at the baseline	–	–	Yes	Yes	–	–	Yes	Yes
Census division indicators	–	–	–	Yes	–	–	–	Yes

Notes: Each column reports results from regressing changes in commuting zone measures of absolute college mobility on decline in manufacturing-to-total employment ratios over the period 1999–2011. The dependent variable is the annual change in college mobility between cohorts born in 1984 and 1993 (N = 616 commuting zones). College mobility is CZ-level average of college attendance of children with parents at the 25th percentile in the national distribution. The college mobility measure comes from Chetty et al. (2014) and is based on the US tax records. The manufacturing decline is instrumented with the commuting zone import exposure from China's imports. The instrument is an employment-weighted average of annualized changes in exposure to Chinese imports within local industries, as detailed in the text. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. The commuting zone baseline controls are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 periods. Census division indicators control for nine regional census divisions. Standard errors are clustered by commuting zone.
$*p < 0.10$
$**p < 0.05$
$***p < 0.01$

Table 4.7: Pretrends: 2SLS Estimates of Manufacturing Decline Effects on High-School Dropout Rates over 1991–2011.

High-School Dropout Rate	(1)	(2)	(3)	(4)	(5)	(6)
Manufacturing decline	-.227***	-.216***	-.172***	-.433***	-.481***	-.418***
	(.031)	(.034)	(.036)	(.142)	(.126)	(.129)
Pretrend controls	–	Yes	Yes	–	Yes	Yes
Baseline controls	–	–	Yes	–	–	Yes
Manufacturing share baseline	–	–	–	Yes	Yes	Yes

Notes: Pretrends. Each column reports results from stacking changes in commuting zone high-school dropout rates and declines in manufacturing-to-total employment ratios over the periods 1991–99 and 1999–2011. The dependent variable is the annual change in the high-school dropout rate (N = 1,444 = 722 commuting zones $\times$ 2 periods). The manufacturing decline is instrumented with the commuting zone import exposure from China's imports. The instrument is an employment-weighted average of annualized changes in exposure to Chinese imports within local industries, as detailed in the text. High-school dropout rate is computed from the US Census for 1990 and 2000, and from the ACS for 2011 as a five-year average. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. All models include a control for time trend. Pretrend controls are annual changes in the high-school dropout rate over 1970–80 and 1980–90 computed from the US Census. Baseline controls include population counts, employment-to-population ratios, and region controls for nine regional census divisions. The commuting zone baseline controls, including the manufacturing share control, are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 periods. Standard errors are clustered by commuting zone.

$*p < 0.10$

$**p < 0.05$

$***p < 0.01$

Table 4.8: Falsification Test: 2SLS Estimates of Manufacturing Decline Effects on High-School Dropout Rates over 1970–2011.

High-School Dropout Rate	1970–80 (1)	1980–90 (2)	1991–99 (3)	1999–2011 (4)
A. Excluding Manufacturing Share Control				
Manufacturing decline 1999–2011	-.314***	.016	-.088**	-.288***
	(.055)	(.058)	(.042)	(.045)
With regional controls	-.044	.048	-.062	-.179***
	(.055)	(.060)	(.046)	(.050)
B. Including Manufacturing Share Control				
Manufacturing decline 1999–2011	-.080	-.213*	-.145	-.487***
	(.141)	(.126)	(.107)	(.138)
With regional controls	.206	-.091	-.111	-.361***
	(.131)	(.126)	(.114)	(.141)

Notes: Falsification test. Each column reports results from a separate specification regressing changes in commuting zone high-school dropout rates in the specified decade and declines in manufacturing-to-total employment ratios over the period 1999–2011. The dependent variable is the annual change in the high-school dropout rate (N = 722 commuting zones over one decade). The manufacturing decline is instrumented with the commuting zone import exposure from China's imports. The instrument is an employment-weighted average of annualized changes in exposure to Chinese imports within local industries, as detailed in the text. High-school dropout rate is computed from the US Census for 1970–2000, and from the ACS for 2011 as a five-year average. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. The commuting zone baseline manufacturing controls are computed in 1999 for the 1999–2011 period. Region controls control for nine regional census divisions. Panels A and B contain no additional controls. Standard errors are clustered by commuting zone.

$*p < 0.10$

$**p < 0.05$

$***p < 0.01$

Table 4.9: Rural vs. Urban: 2SLS Estimates of Trade Exposure Effects on High-School Dropout Rate 1991–2011 and College Mobility 1999–2011.

Rural vs. Urban	2SLS	
	(1)	(2)
	A. High-School Dropout Rate	
Commuting zone manufacturing decline	-.224***	-.407***
	(.048)	(.129)
Interaction: manufacturing decline × rural	-.027	.029
	(.061)	(.067)
Baseline manufacturing emp. share	–	Yes
Other baseline controls	–	Yes
	B. College Mobility	
Commuting zone manufacturing decline	.566***	.100
	(.142)	(.261)
Interaction: manufacturing decline × rural	-.167	-.21
	(.235)	(.235)
Baseline manufacturing emp. share	–	Yes
Other baseline controls	–	Yes

Notes: Rural vs. Urban. In Panel A, each column reports results from stacking the logarithms of changes in commuting zone high-school dropout rates and declines in manufacturing-to-total employment ratios over the periods 1991–99 and 1999–2011. The dependent variable is the annual change in the high-school dropout rate (N = 1,444 = 722 commuting zones × 2 periods). High-school dropout rate is computed from the US Census for 1990 and 2000, and from the ACS for 2011 as a five-year average. In Panel B, each column reports results from regressing changes in commuting zone measures of absolute college mobility and declines in manufacturing-to-total employment ratios over the periods 1991–99 and 1999–2011. The dependent variable is the annual change in college mobility between cohorts born in 1984 and 1993 (N = 616 commuting zones). College mobility is CZ-level average of college attendance of children with parents at the 25th percentile in the national distribution. The college mobility measure comes from Chetty et al. (2014) and is based on the US tax records. In both Panels A and B, manufacturing decline is instrumented with the commuting zone import exposure from China's imports. The instrument is an employment-weighted average of annualized changes in exposure to Chinese imports within local industries, as detailed in the text. Both panels include interaction terms with US Census rural area indicator as in text. The commuting zone baseline controls are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 period. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. The other baseline controls include population counts, employment-to-population ratios, and region controls for nine regional census divisions. All models include a time trend. Standard errors are clustered by commuting zone.

$*p < 0.10$

$**p < 0.05$

$***p < 0.01$

Table 4.10: Geographical Correlates of the Intergenerational Effects of Manufacturing Decline.

	2SLS	
Interaction term	Main effect	Interaction
Segregation and Race		
Fraction Black	-0.213* (0.112)	-0.820** (0.363)
Income Segregation	-0.311** (0.131)	-2.239* (1.177)
Segregation of Affluence (>p75)	-0.308** (0.131)	-2.154** (1.063)
Fraction with Commute < 15 Mins	-0.596*** (0.143)	0.851*** (0.243)
Income Inequality		
Household Income per Capita	-0.190 (0.242)	-0.000 (0.000)
Gini coefficient	-0.281 (0.192)	-0.214 (0.331)
Fraction Middle Class (between p25 and p75)	-0.573** (0.280)	0.472 (0.427)
K-12 Education		
School Expenditure per Student	-0.465** (0.200)	0.021 (0.034)
Student Teacher Ratio	-0.053 (0.251)	-0.021 (0.015)
Test Score Percentile (Income adjusted)	-0.300*** (0.110)	-0.002 (0.004)
College		
Number of Colleges per Capita	-0.577*** (0.156)	5.757*** (1.823)
College Tuition	-0.458*** (0.164)	-0.000 (0.000)
College Graduation Rate (Income Adjusted)	-0.483*** (0.161)	-0.000 (0.000)
Social Capital		
Social Capital Index	-0.362*** (0.126)	0.038 (0.029)
Fraction Religious	-0.427** (0.167)	0.107 (0.292)
Violent Crime Rate	-0.317** (0.139)	-42.456 (36.070)
Local Labor Market		
Teenage (14-16) Labor Force Participation	-0.475*** (0.162)	39.034* (23.495)

Notes: Each column reports results from stacking the logarithms of changes in commuting zone high-school dropout rates and declines in manufacturing-to-total employment ratios over the periods 1991–99 and 1999–2011, and including an interaction term and main effect for the indicated variable (N = 1,444 = 722 commuting zones × 2 periods). The baseline controls include manufacturing share of employment, population counts, employment-to-population ratios, and region controls for nine regional census divisions. All models include a time trend. Standard errors are clustered by commuting zone. The variables are detailed in Tables D.4 and D.5. Further details are provided in text.

$*p < 0.10$

$**p < 0.05$

$***p < 0.01$

Appendix A

Appendix to Chapter 1

A.1 Winners-Losers: Supplementary Figures and Tables

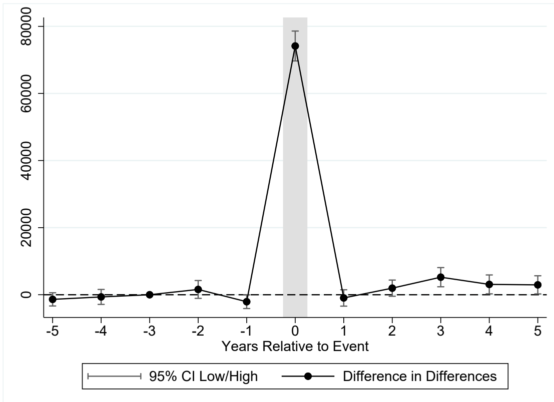

(a) Granted.

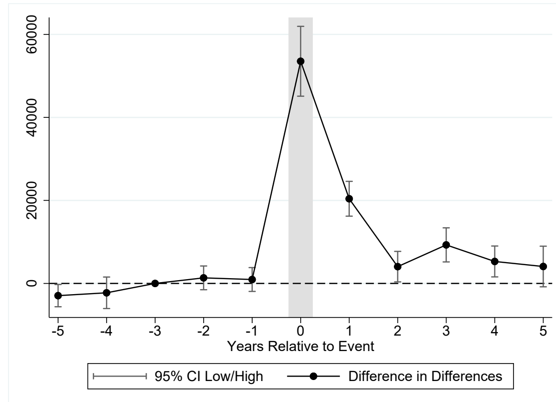

(b) Received.

Figure A-1: The First Stage: The Effect of Winning a Subsidy on Granted and Received Subsidies.

Notes: Event-study estimates from Equation 1.1. Panel (a): The outcomes are (a) any subsidy granted and (b) received, measured from the Finnish Statistics on Business Subsidies. Event time $\tau = 0$ refers to the application year. Back to Section 1.5.

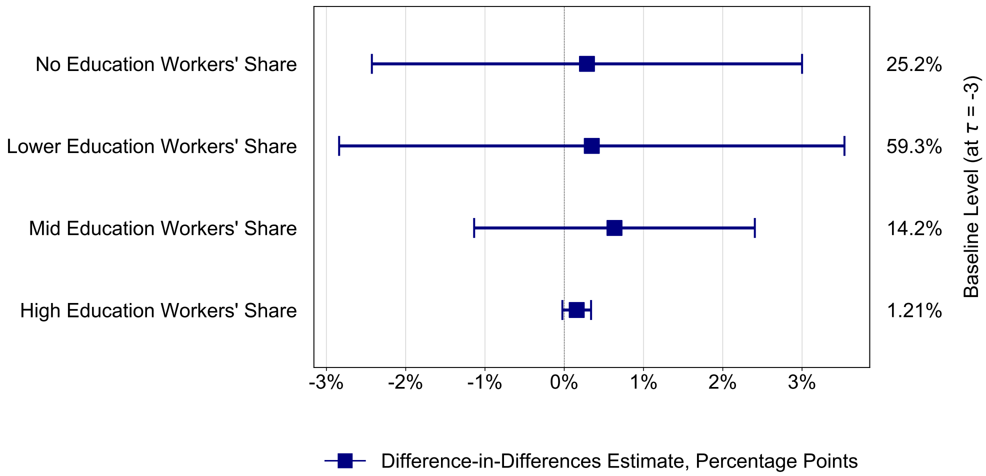

Figure A-2: Skill Effects: Education Groups.

Notes: Difference-in-differences estimates from Equation 1.2. The right-hand side reports means at $\tau = -3$. The data are from Finnish educational registers. Education groups are defined in Appendix A.5. Back to Section 1.5.

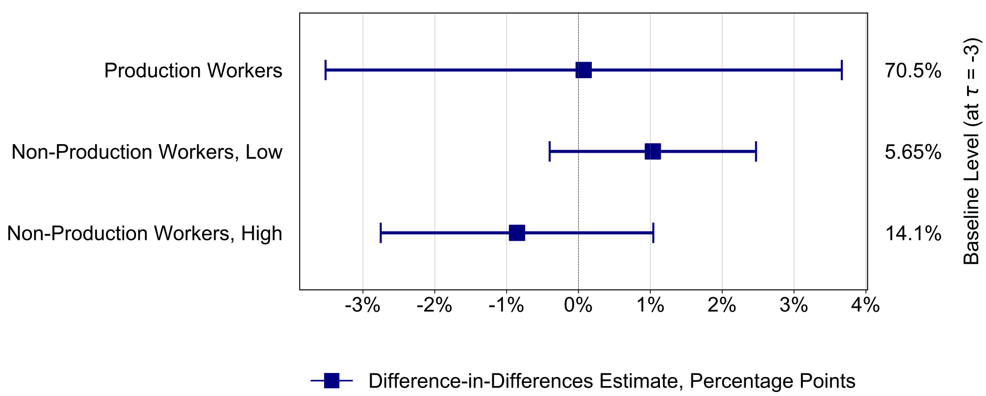

Figure A-3: Skill Effects: Occupation Groups.

Notes: Difference-in-differences estimates from Equation 1.2. The right-hand side reports means at $\tau = -3$. The data are from the Finnish occupation registers. Occupation groups are defined in Appendix A.5. The shares do not sum to 100% because some workers do not have occupational info, i.e., the denominator includes all workers in the firm. Back to Section 1.5.

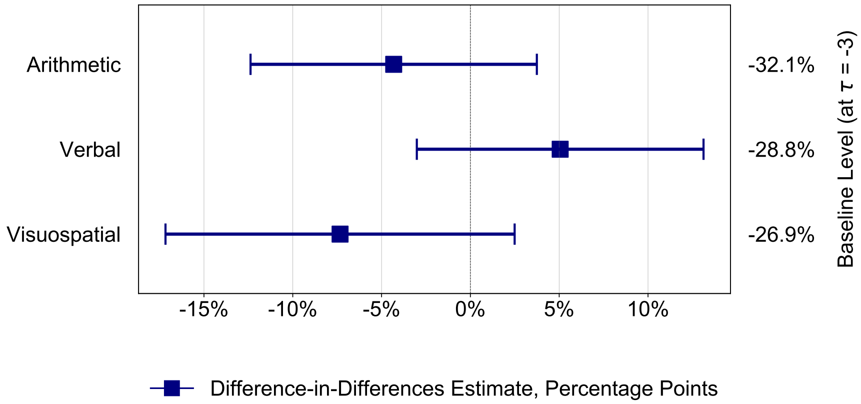

Figure A-4: Skill Effects: Cognitive Performance.

Notes: Difference-in-differences estimates from Equation 1.2. The right-hand side reports means at $\tau = -3$. The estimates are in percentages of standard deviations. The data are from the Finnish Defence Forces. Back to Section 1.5.

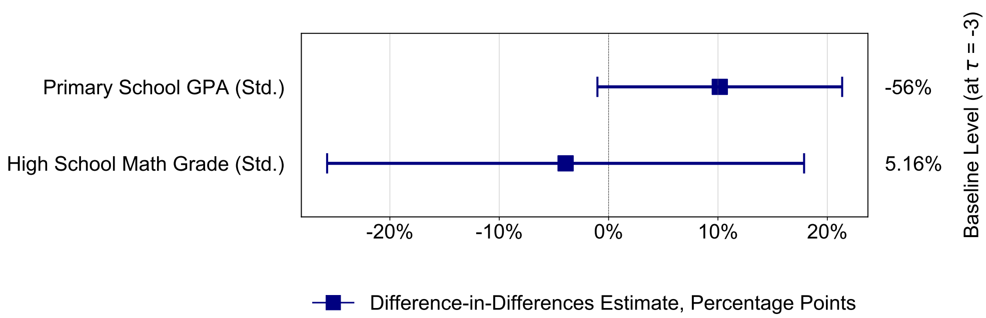

Figure A-5: Skill Effects: School Performance.

Notes: Difference-in-differences estimates from Equation 1.2. The right-hand side reports means at $\tau = -3$. The estimates are in percentages of standard deviations. The data are from the Secondary Education Application Register and the Finnish Matriculation Examination Board Register. Back to Section 1.5.

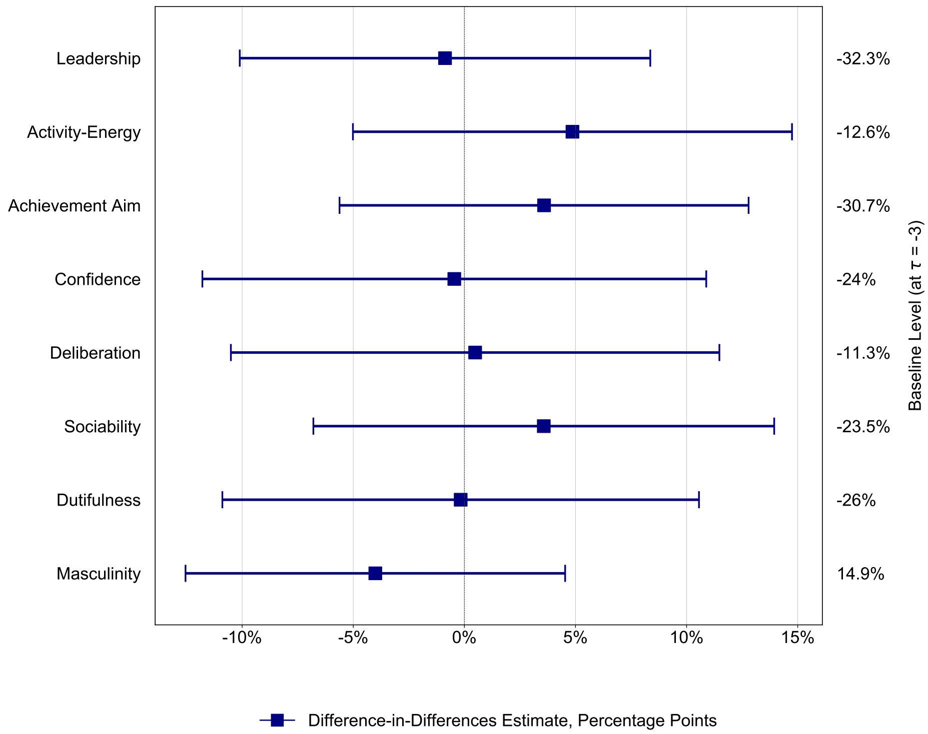

Figure A-6: Skill Effects: Personality.

Notes: Difference-in-differences estimates from Equation 1.2. The right-hand side reports means at $\tau = -3$. The estimates are in percentages of standard deviations. The data are from the Finnish Defence Forces. Back to Section 1.5.

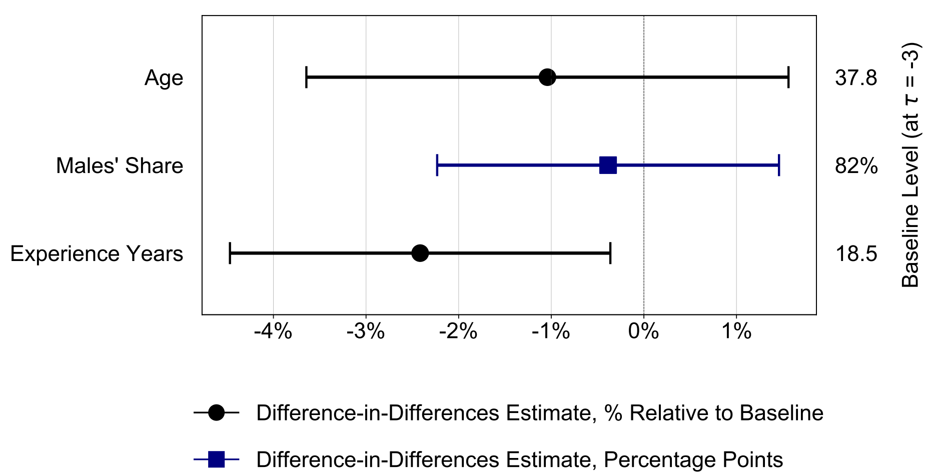

Figure A-7: Skill Effects: Demographics.

Notes: Difference-in-differences estimates from Equation 1.2. The right-hand side reports means at $\tau = -3$. The data are from the Finnish worker and population registers. Back to Section 1.5.

Panel A: Above Median.

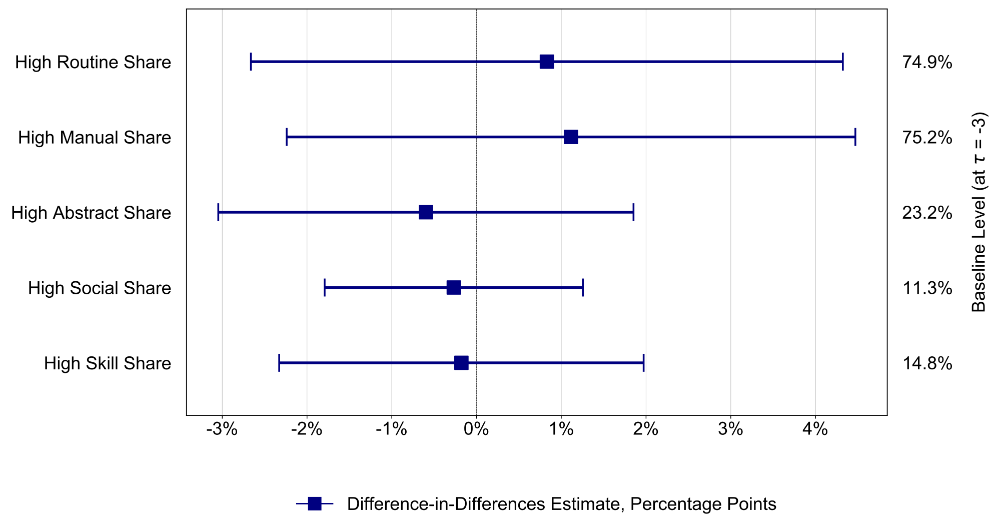

Panel B: Below Median.

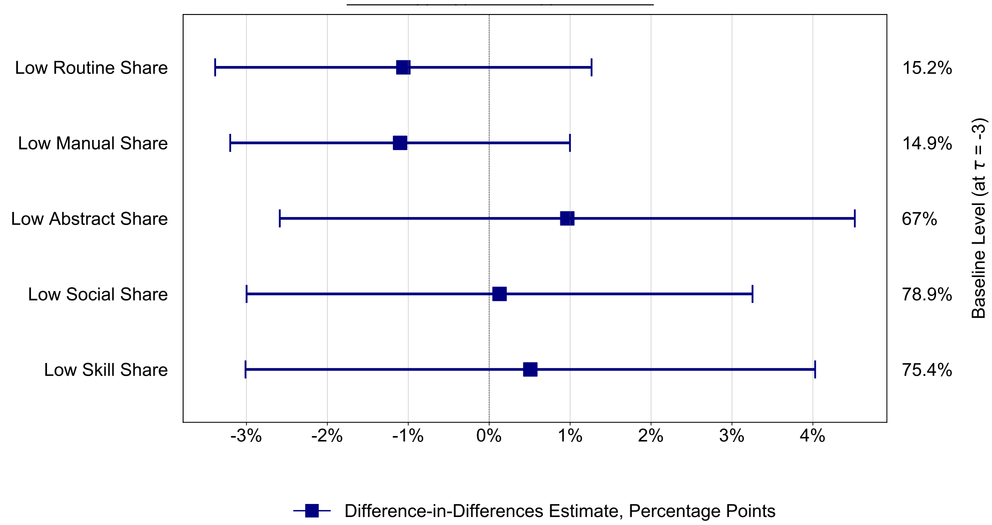

Figure A-8: Skill Effects: Tasks.

Notes: Difference-in-differences estimates from Equation 1.2. The right-hand side reports means at $\tau = -3$. Median refers to the median task intensity in the Finnish labor force. For example, the first row indicates that 74.9% of workers in our sample firms are in an occupation times industry cell that is above the median in routine task content. The treatment group increases the share of these workers by a statistically insignificant 1% compared to the control group. The shares do not sum to 100% because some workers do not have occupational info (the denominator includes all workers in the firm). The data are from the Finnish occupation registers and the European Working Conditions Survey (EWCS). Back to Section 1.5.

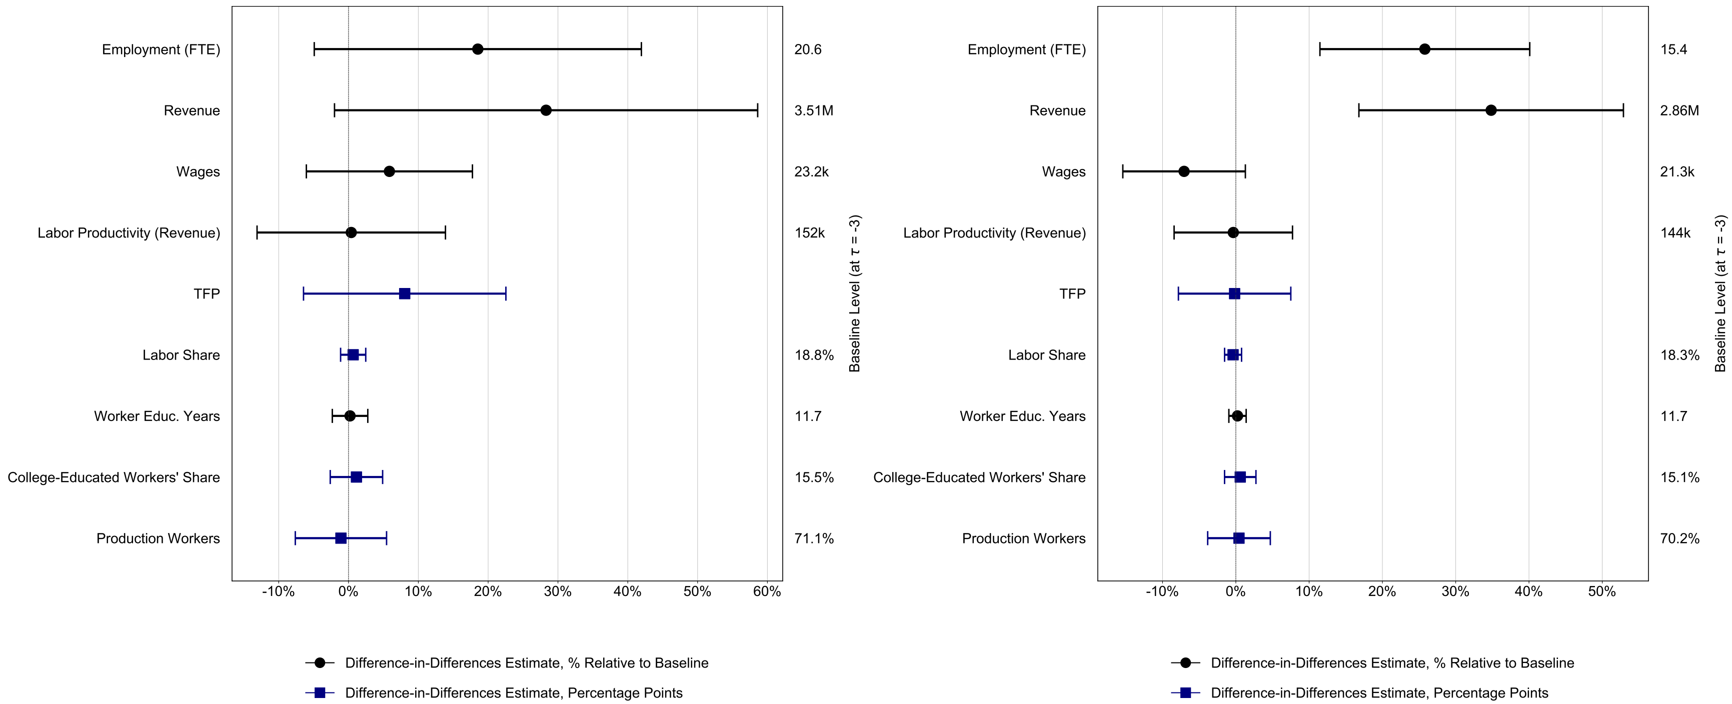

Figure A-9: Automated (left) vs. Non-Automated (right) Technologies from Text Data.

Notes: Difference-in-differences estimates from Equation 1.2. Automated vs. non-automated technologies are measured from text data as described in Section 1.3 and Appendix A.5. Automated (N): Treatment 678, Control 30. Non-Automated (N): Treatment 1207, Control 116. Back to Section 1.7.

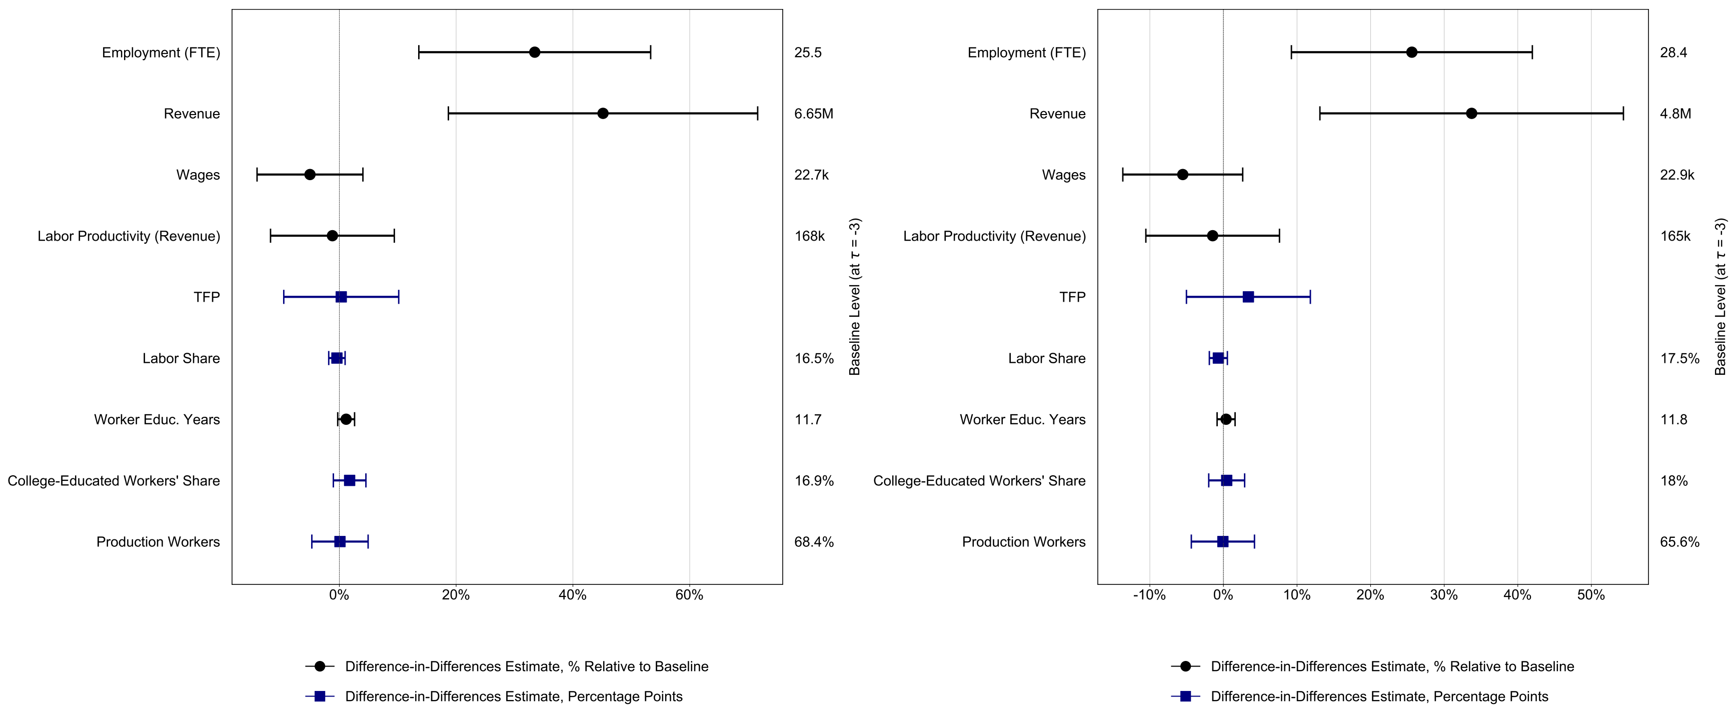

Figure A-10: Automated (left) vs. Non-Automated (right) Technologies from Customs Data.

Notes: Difference-in-differences estimates from Equation 1.2. Automated vs. non-automated technologies are measured from customs data as described in Section 1.3 and Appendix A.5. A project is classified as automated if over 50% of the imported machinery are automated technologies. A project is classified as non-automated if over 50% of the imported machinery are non-automated technologies. Automated (N): Treatment 220, Control 146. Non-Automated (N): Treatment 319, Control 146. Back to Section 1.7.

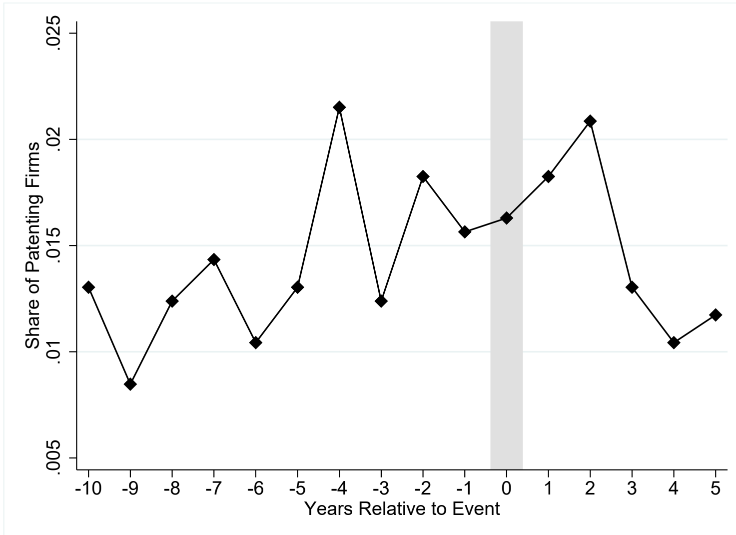

Figure A-11: Patents: Share of Patenting Firms.

Notes: The share of patenting firms by year among subsidy applicant firms. Patent information comes from the Finnish Patent Database. Event time $\tau = 0$ refers to the subsidy application year. Back to Section 1.6.2.

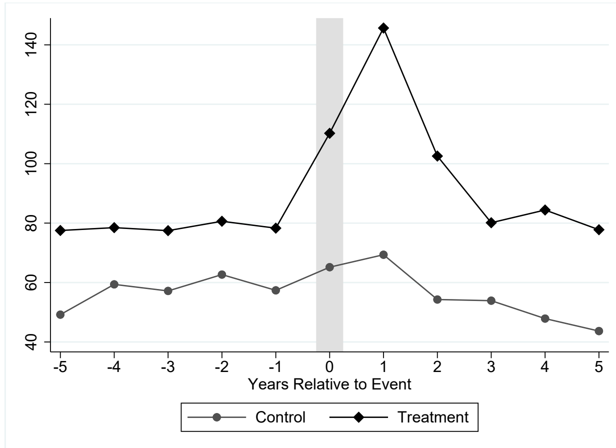

(a) Machinery Investment.

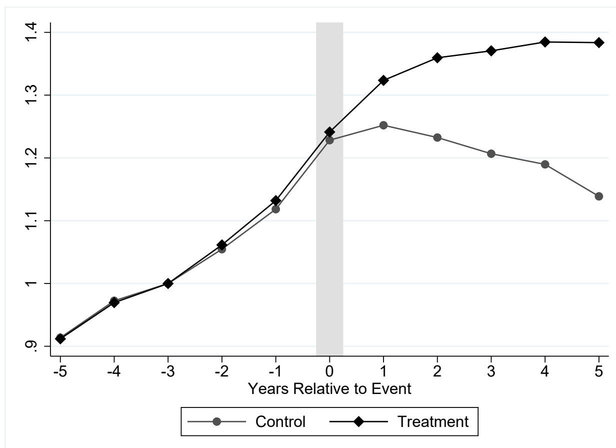

(b) Employment.

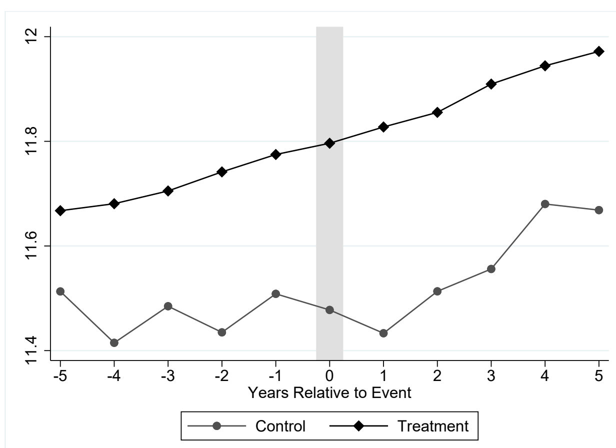

(c) Education Years.

Figure A-12: Raw Means: Machinery Investment, Employment, and Education.

Notes: Means over time for the main treatment and control groups (winners vs. losers). Machinery investment in EUR, employment in % relative to $\tau = -3$, and education in years. The patterns in the main control group are similar to the patterns in a matched non-applicant control group as shown by Figure A-29. Back to Section 1.7.

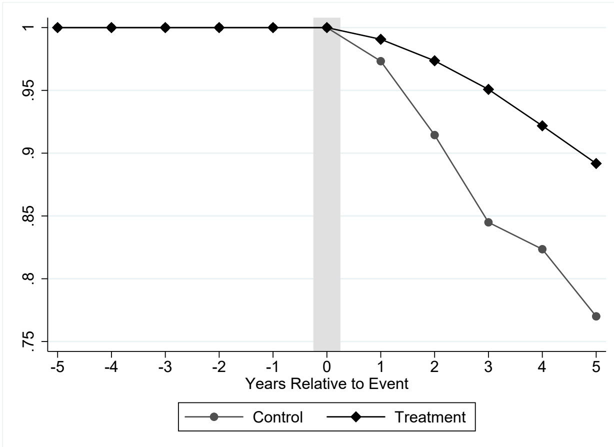

(a) Firm Survival Based on the Firm Register.

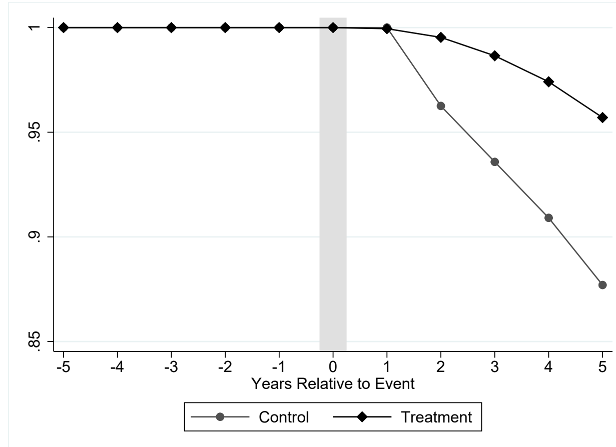

(b) Firm Survival Based on Worker Flows.

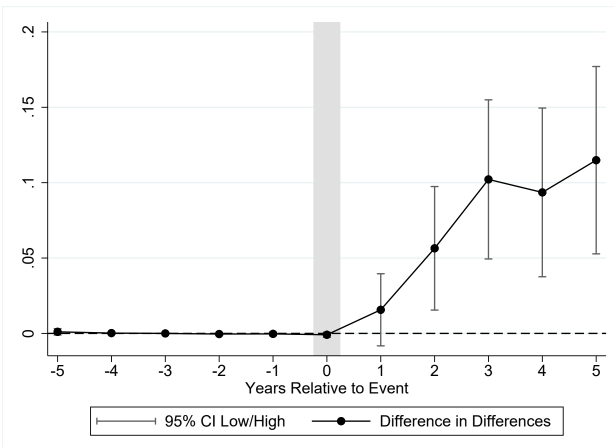

(c) Firm Survival Based on the Firm Register.

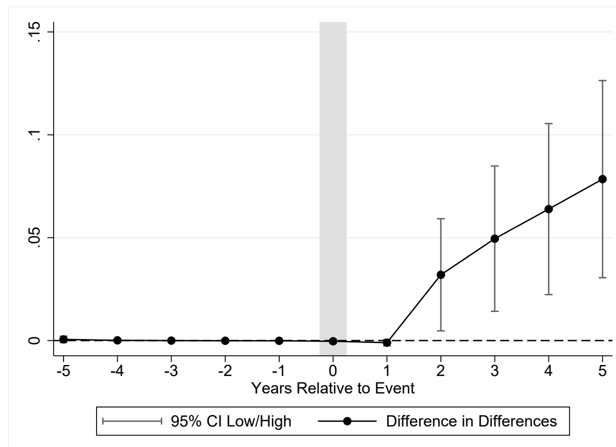

(d) Firm Survival Based on Worker Flows.

Figure A-13: Firm Survival Effects.

Notes: Group means and event-study estimates from Equation 1.1. **Panels (a, c):** Survival is measured from whether the firm ID exists in the firm register. **Panels (b, d):** Survival is extended to include mergers and acquisitions (and other cases the firm ID changes), where at least 50% of workers continue under the same firm ID. The main estimates are reported for a balanced sample over the 5-year window. The estimates are robust to a non-balanced sample, shown in Table A.14. Back to Section 1.7.

Table A.1: Summary Statistics: Benchmarking to All Manufacturing.

	Subsidy Sample				**Finnish Manufacturing**			
Variable	Mean	p10	Median	p90	Mean	p10	Median	p90
Revenue (EUR M)	2.66	1.96	2.56	3.77	2.03	1.89	2.01	2.27
Employment	16.25	12.76	16.11	19.07	12.35	11.63	12.52	13.06
Wages (EUR K)	26.25	19.98	25.88	32.83	26.95	21.13	26.94	32.16
Labor Productivity (EUR K)	150.30	131.20	147.80	171.08	140.55	125.82	142.72	152.31
Profit Margin (%)	5.55	3.10	5.56	7.63	4.47	2.94	4.56	5.84
Employment Change (%, Five Year)	57.72	40.70	50.81	84.15	48.11	34.52	44.24	82.17
Revenue Change (%, Five Year)	74.62	44.66	74.76	96.19	59.87	30.25	54.83	101.80
Subsidy Applied (EUR K)	110.48	86.74	107.61	149.45	4.80	3.38	4.68	6.20
Subsidy Granted (EUR K)	79.53	49.82	78.51	109.14	2.58	2.13	2.62	3.27
Educ. Years	11.79	11.57	11.77	12.07	11.64	11.49	11.60	11.84
College Share (%)	15.36	13.38	15.37	17.94	14.56	13.33	14.78	15.45
Production Worker Share (%)	70.70	66.37	69.99	74.87	69.33	66.76	69.12	72.67
Number of Observations	2031				260,220			
Number of Unique Firms	2031				18,501			
Number of Years	16				16			

Notes: Manufacturing firms include all firms that satisfy the subsidy sample's balance-sheet-based restrictions and have over two full-time employees. The subsidy sample is measured at event-time $\tau = -1$. Manufacturing means are measured for each firm in a given year and collapsed to a year-level mean for all manufacturing. These year-level means are averaged over 1994–2018. The median and the percentiles are at the year level. Subsidy applied, subsidy granted, college share, and production worker share are not winsorized, but all other outcomes are (at top and bottom 5% level). Back to Section 1.4.1.

Table A.2: Summary Statistics: Text Matching using Cosine Similarity.

	Treatment Group		Control Group		Both		
Variable	Mean	Std. Dev.	Mean	Std. Dev.	10p	Median	90p
Revenue (EUR M)	2.26	4.44	1.68	3.85	0.13	0.72	4.68
Employment	15.77	26.04	11.15	24.65	1.10	5.90	27.40
Wages (EUR K)	21.24	8.15	19.28	10.29	6.73	21.27	29.23
Subsidy Applied (EUR K)	110.02	128.33	64.64	105.44	4.60	38.35	241.32
Subsidy Granted (EUR K)	78.31	99.14	0.00	0.00	0.00	0.34	124.65
Educ. Years	11.67	0.98	11.42	1.04	10.50	11.63	12.50
College Share (%)	15.18	16.75	11.05	16.30	0.00	10.30	33.33
Production Worker Share (%)	70.62	22.17	72.65	27.18	40.00	75.00	100.00
Observations	1508		1508		3016		

Notes: All variables measured at $\tau = -3$. Back to Section 1.4.3.

Table A.3: Firm-Level Effects: Different Text Matching Versions.

Panel A: Coarsened Exact Matching (CEM).

	(1) Machine Inv. (EUR K)	(2) Employment	(3) Revenue	(4) Educ. Years	(5) College Share	(6) Production Worker Share
Treatment	93.10***	0.242***	0.313**	-0.0480	-0.000144	-0.00883
	(19.93)	(0.0712)	(0.0956)	(0.0661)	(0.0105)	(0.0207)
Observations	1256	1256	1256	1160	1160	1161

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel B: Inverse Probability Weighting (IPW).

	(1) Machine Inv. (EUR K)	(2) Employment	(3) Revenue	(4) Educ. Years	(5) College Share	(6) Production Worker Share
Treatment	159.6***	0.359***	0.458***	-0.0441	0.00547	-0.0276
	(22.81)	(0.0911)	(0.117)	(0.0848)	(0.0162)	(0.0300)
Observations	1812	1812	1812	1676	1676	1692

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel C: Cosine Similarity.

	(1) Machine Inv. (EUR K)	(2) Employment	(3) Revenue	(4) Educ. Years	(5) College Share	(6) Production Worker Share
Treatment	103.9***	0.169***	0.195***	0.0133	-0.00224	-0.00769
	(14.90)	(0.0249)	(0.0335)	(0.0219)	(0.00542)	(0.00896)
Observations	3016	3016	3016	2678	2678	2678

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2 with different text matching versions. Back to Section 1.5.

Table A.4: Firm-Level Effects: Different Controls.

	(1) Machine Inv. (EUR K)	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Share	(7) Educ. Years	(8) College Share	(9) Prod. Work. Share
No Controls	140.9***	0.185**	0.261***	-0.0553	-0.00559	-0.00242	0.0225	0.00480	-0.0235
	(25.76)	(0.0606)	(0.0770)	(0.0353)	(0.0341)	(0.00492)	(0.0599)	(0.00927)	(0.0359)
Controls 1	132.4***	0.219***	0.302***	-0.0499	-0.00379	-0.00247	0.0252	0.00587	-0.0263
	(26.17)	(0.0615)	(0.0779)	(0.0356)	(0.0351)	(0.00496)	(0.0611)	(0.00936)	(0.0357)
Controls 2	114.8***	0.232***	0.314***	-0.0481	-0.00516	-0.00202	0.0246	0.00557	-0.0256
	(23.99)	(0.0614)	(0.0779)	(0.0355)	(0.0350)	(0.00496)	(0.0611)	(0.00935)	(0.0357)
Controls 3	105.0***	0.249***	0.327***	-0.0385	-0.00670	-0.000862	0.0252	0.00572	-0.0255
	(23.96)	(0.0609)	(0.0773)	(0.0350)	(0.0349)	(0.00490)	(0.0612)	(0.00942)	(0.0363)
Controls 4	41.02	0.210***	0.284***	-0.0344	-0.00658	-0.000101	0.0247	0.00509	-0.0268
	(22.92)	(0.0607)	(0.0770)	(0.0351)	(0.0350)	(0.00493)	(0.0614)	(0.00946)	(0.0363)
Controls 5	36.43	0.221***	0.299***	-0.0319	-0.00474	-0.000143	0.0168	0.00482	-0.0275
	(22.65)	(0.0613)	(0.0776)	(0.0352)	(0.0350)	(0.00494)	(0.0619)	(0.00951)	(0.0366)
Observations	2031	2031	2031	1952	2031	2031	1884	1884	821

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2 with different controls.

Controls 1: industry (2-digit).

Controls 2: industry (2-digit), employment (at the base year).

Controls 3: industry (2-digit), employment (at the base year), ELY Center indicators.

Controls 4: industry (2-digit), employment (at the base year), ELY Center indicators, applied subsidy amount.

Controls 5: industry (2-digit), employment (at the base year), ELY Center indicators, applied subsidy amount, text category indicators.

Back to Section 1.5.

Table A.5: Continuous Treatment Estimates Controlling for the Subsidies Applied.

	(1) Machine Inv. (EUR K)		(2) Employment		(3) Revenue	
Granted Subsidy	0.589***	0.613***	0.129**	0.140**	1.546	2.074*
	(0.153)	(0.163)	(0.0464)	(0.0500)	(0.960)	(1.038)
Applied Subsidy	✓	✓	✓	✓	✓	✓
Propensity Score		✓		✓		✓
Observations	2031	1812	2031	1812	2031	1812

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2. Treatment is the received subsidy amount in EUR. Treatment is scaled to EUR 10K for employment. Applied subsidy is the applied subsidy amount in EUR. Machinery investment is the sum over $\tau \in [0, 2]$. Other outcomes are averages over $\tau \in [2, 5]$. Back to Section 1.5.

Table A.6: Product: Matched Sample Summary Statistics.

	Treatment Group		Control Group		Both		
Variable	Mean	Std. Dev.	Mean	Std. Dev.	10p	Median	90p
Revenue (EUR M)	3.65	33.13	6.42	23.70	0.16	1.02	7.52
Employment	18.65	55.00	30.54	87.85	1.50	8.50	43.30
Wages (EUR K)	22.23	8.27	22.74	8.68	12.95	23.05	31.17
Subsidy Applied (EUR K)	111.99	128.51	3.13	23.43	0.00	4.03	182.69
Subsidy Granted (EUR K)	83.63	104.87	1.87	13.83	0.00	2.94	131.11
Educ. Years	11.71	1.00	11.62	1.04	10.50	11.70	12.67
College Share (%)	15.38	16.94	16.05	18.42	0.00	12.50	35.23
Production Worker Share (%)	70.81	21.92	67.97	24.67	37.50	72.34	100.00
Observations	1023		1023		2046		

Notes: All variables measured at $\tau = -3$. The treatment group is subsidy-winning firms that described product-type technological advances in their application text. The matched control group is searched from all non-applicant firms with balance sheet data. In this table, the subsidy applied and granted refer to all recorded subsidies; the matched control group does not apply or receive ELY Center subsidies. Back to Section 1.6.2.

Table A.7: Process: Matched Sample Summary Statistics.

	Treatment Group		Control Group		Both		
Variable	Mean	Std. Dev.	Mean	Std. Dev.	10p	Median	90p
Revenue (EUR M)	3.06	6.22	3.18	5.29	0.16	1.02	8.14
Employment	21.61	38.00	21.85	37.54	1.30	8.80	46.50
Wages (EUR K)	23.67	8.34	23.95	8.71	14.68	24.19	33.67
Subsidy Applied (EUR K)	77.50	95.55	13.12	59.09	0.00	4.16	141.99
Subsidy Granted (EUR K)	52.94	72.32	8.22	35.42	0.00	3.49	90.19
Educ. Years	11.57	0.95	11.53	0.93	10.50	11.68	12.52
College Share (%)	14.45	15.99	14.50	16.75	0.00	12.50	30.60
Production Worker Share (%)	69.48	20.42	70.32	22.70	50.00	71.43	100.00
Observations	99		99		198		

Notes: All variables measured at $\tau = -3$. The treatment group is subsidy-winning firms that described process-type technological advances in their application text. The matched control group is searched from all non-applicant firms with balance sheet data. In this table, the subsidy applied and granted refer to all recorded subsidies; the matched control group does not apply or receive ELY Center subsidies. Back to Section 1.6.2.

Table A.8: The Effects by Technology Categories Measured from Text Data.

Panel A: Investment, Employment, Wages, and Firm Performance.

	(1) Machine Inv. (EUR K)	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity
Product	142.7***	0.210***	0.262***	-0.00270	0.0222
	(9.964)	(0.0235)	(0.0320)	(0.0122)	(0.0154)
Process	77.66***	0.0905	0.0783	-0.00154	-0.0515
	(22.95)	(0.0779)	(0.0759)	(0.0324)	(0.0483)
N, Product	2046	2046	2046	1963	2046
N, Process	198	198	198	192	198

Panel B: Skill Composition and The Labor Share.

	(1) Labor Share	(2) Educ. Years	(3) College Share	(4) Production Worker Share
Product	-0.00474	0.0227	0.00691	-0.0110
	(0.00264)	(0.0269)	(0.00422)	(0.00742)
Process	0.00583	0.137	0.00497	0.0101
	(0.00765)	(0.0809)	(0.0135)	(0.0211)
N, Product	2046	1905	1905	1921
N, Process	198	186	186	186

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2. Product (the extensive margin) refers to technology projects that aim to produce a new type of output. Process (the intensive margin) refers to technology projects that aim to produce the same type of output with new technologies. **Panel A:** Column 1 is in EUR K. Columns 2, 3, 4, and 5 are relative changes, e.g., 0.20 would refer to a 20% increase. **Panel B:** Columns 1, 3, and 4 (shares) are in percentage points. Column 2 (education) is in years. We use coarsened exact matching (CEM) to construct the control group. N refers to the number of matched observations. For machine investment, the post-period outcome is the sum of investment between $\tau \in [0, 2]$ and for other outcomes, the average of $\tau \in [2, 5]$. Back to Section 1.6.2.

Table A.9: The Effects by Technology Categories Measured from Survey Data.

Panel A: Employment, Wages, and Firm Performance.

	(1) Machine Inv. (EUR K)	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity
Product	311.5***	0.235**	0.364***	-0.00137	0.154*
	(62.75)	(0.0812)	(0.101)	(0.0296)	(0.0620)
Process	-	-	-	-	-
	-	-	-	-	-
N, Product	164	164	164	164	164
N, Process	6	6	6	6	6

Panel B: Skill Composition and the Labor Share.

	(1) Labor Share	(2) Educ. Years	(3) College Share	(4) Production Worker Share
Product	-0.0169*	0.0758	0.00812	-0.00478
	(0.00737)	(0.0679)	(0.0107)	(0.0184)
Process	-	-	-	-
	-	-	-	-
N, Product	164	163	163	163
N, Process	6	6	6	6

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2. The technology categories are measured from the European Community Innovation Survey (CIS). Product (the extensive margin) refers to technology projects that aim to produce a new type of output. Process (the intensive margin) refers to technology projects that aim to produce the same type of output with new technologies. The process sample is too small to perform estimation (denoted by –). **Panel A:** Column 1 is in EUR K. Columns 2, 3, 4, and 5 are relative changes, e.g., 0.20 would refer to a 20% increase. **Panel B:** Columns 1, 3, and 4 (shares) are in percentage points. Column 2 (education years) is in years. We use coarsened exact matching (CEM) 1:1. N refers to matched observations. Machine investment is the sum over $\tau \in [0, 2]$, other outcomes are averages over $\tau \in [2, 5]$. Back to Section 1.6.2.

Table A.10: The Effects by Context Measured from the Rauch Index.

Panel A: Employment, Wages, and Firm Performance.

	(1) Machine Inv. (EUR K)	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity
Specialized	147.9***	0.188***	0.216***	-0.00748	0.00401
	(8.141)	(0.0213)	(0.0272)	(0.0113)	(0.0134)
Non-Specialized	86.61*	0.132	0.171	0.0334	0.0122
	(42.06)	(0.0965)	(0.114)	(0.0386)	(0.0612)
N, Specialized	2704	2704	2704	2606	2704
N, Non-Specialized	248	248	248	242	248

Panel B: Skill Composition and the Labor Share.

	(1) Labor Share	(2) Educ. Years	(3) College Share	(4) Production Worker Share
Specialized	-0.00184	0.0247	0.00281	-0.00350
	(0.00219)	(0.0218)	(0.00361)	(0.00637)
Non-Specialized	-0.00149	-0.00469	-0.00735	0.0399
	(0.00988)	(0.107)	(0.0192)	(0.0251)
N, Specialized	2704	2539	2539	2584
N, Non-Specialized	248	236	236	239

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2. **Panel A:** Column 1 is in EUR K. Columns 2, 3, 4, and 5 are relative changes, e.g., 0.20 would refer to a 20% increase. **Panel B:** Columns 1, 3, and 4 (shares) are in percentage points. Column 2 (education years) is in years. N refers to matched observations. We use coarsened exact matching 1:1 (CEM). Machine investment is the sum over $\tau \in [0, 2]$. Other outcomes are averages over $\tau \in [2, 5]$. Details in the main text. Back to Section 1.6.3.

Table A.11: Technology Categories from Text Data vs. Rauch Index.

Class	Product	Process	Total
High Rauch Index	1019	89	1108
Low Rauch Index	98	15	113
Total	1117	104	1221

Notes: This 2x2 table reports the number of firms in the text categories and Rauch Index combinations. Product refers to technology projects that aim to produce a new type of output. Process refers to technology projects that aim to produce the same type of output with new technologies. High Rauch Index refers to specialized industries, Low Rauch Index refers to non-specialized industries. Back to Section 1.6.3.

Table A.12: The Effects by Firm Size.

Panel A: Large Firms.

	(1) Machine Inv. (EUR K)		(2) Employment		(3) Revenue		(4) Productivity		(5) Labor Share		(6) College Share	
Treatment	83.88	68.81	0.305***	0.309**	0.264	0.494**	-0.136	0.0236	0.0133	-0.00853	-0.00893	-0.0159
	(69.06)	(88.72)	(0.0722)	(0.104)	(0.137)	(0.160)	(0.0834)	(0.0967)	(0.00981)	(0.0111)	(0.0167)	(0.0201)
Propensity Score		✓		✓		✓		✓		✓		✓
Observations	676	609	676	609	676	609	676	609	676	609	675	608

Panel B: Medium-Sized Firms.

	(1) Machine Inv. (EUR K)		(2) Employment		(3) Revenue		(4) Productivity		(5) Labor Share		(6) College Share	
Treatment	76.82*	87.38*	0.296***	0.280*	0.467***	0.399**	0.0707	0.0193	-0.0104	-0.0124	0.0185	0.0200
	(33.67)	(41.97)	(0.0858)	(0.113)	(0.114)	(0.150)	(0.0551)	(0.0718)	(0.00856)	(0.00969)	(0.0162)	(0.0213)
Propensity Score		✓		✓		✓		✓		✓		✓
Observations	685	603	685	603	685	603	685	603	685	603	683	601

Panel C: Small Firms.

	(1) Machine Inv. (EUR K)		(2) Employment		(3) Revenue		(4) Productivity		(5) Labor Share		(6) College Share	
Treatment	31.99*	28.23	0.330**	0.373**	0.355**	0.370*	-0.0410	-0.0956	0.00216	0.0162	0.00334	0.00373
	(13.48)	(18.09)	(0.103)	(0.121)	(0.125)	(0.148)	(0.0526)	(0.0615)	(0.00781)	(0.00927)	(0.0158)	(0.0192)
Propensity Score		✓		✓		✓		✓		✓		✓
Observations	670	600	670	600	670	600	670	600	670	600	526	467

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Difference-in-differences estimates from Equation 1.2. Large Firms (FTE > 13.3; Median 25.8, Mean 41.7), Medium-Sized Firms (FTE >= 4.6 & FTE <= 13.3; Median 7.9, Mean 8.2), Small Firms (FTE < 4.6; Median 2.3, Mean 2.3). Back to Section 1.6.3.

Table A.13: Credit Constraints: Robustness Checks.

Panel A: Effects by Average Financial Costs.

	(1) Machine Inv. (EUR K)		(2) Employment		(3) Revenue	
	High Costs	Low Costs	High Costs	Low Costs	High Costs	Low Costs
Treatment	113.3***	91.37***	0.276***	0.194*	0.313**	0.326**
	(25.88)	(23.11)	(0.0793)	(0.0959)	(0.105)	(0.114)
Observations	1016	1015	1016	1015	1016	1015

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel B: Effects by Relative Debt.

	(1) Machine Inv. (EUR K)		(2) Employment		(3) Revenue	
	High Debt	Low Debt	High Debt	Low Debt	High Debt	Low Debt
Treatment	121.9***	78.93***	0.0676	0.384***	0.151	0.486***
	(25.97)	(23.58)	(0.0965)	(0.0687)	(0.125)	(0.0820)
Observations	1016	1015	1016	1015	1016	1015

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel C: Controlling for Credit Constraint Measures.

	(1) Machine Inv. (EUR K)			(2) Employment			(3) Revenue		
Treatment	107.9***	107.8***	108.3***	0.232***	0.242***	0.232***	0.314***	0.342***	0.314***
	(17.53)	(17.54)	(17.59)	(0.0614)	(0.0597)	(0.0614)	(0.0779)	(0.0721)	(0.0778)
Relative Debt		✓			✓			✓	
Average Financial Costs			✓			✓			✓
Observations	2031	2031	2031	2031	2031	2031	2031	2031	2031

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The sample is the main analysis sample (subsidies design). Estimated effects on selected outcomes by the cost of capital (Panel A) and debt level (Panel B), and with credit-constraint controls (Panel C). We measure baseline levels at $\tau = -3$. Average financial costs are financial expenses divided by non-current liabilities. Relative debt is the sum of current liabilities, non-current liabilities, and obligatory reserves divided by revenue. We divide the sample into two groups by whether the firms' average financial costs (Panel A) or relative debt (Panel B) are below or above the median in the sample. Panel C controls directly for the baseline value.

Table A.14: Robustness to a Non-Balanced Sample: Firm-Level Effects Allowing for Firm Exit.

Panel A: Investment, Employment, Wages, and Firm Performance.

	(1)		(2)		(3)		(4)		(5)	
	Machine Inv. (EUR K)		Employment		Revenue		Wages		Profit Margin	
Treatment	93.16***	103.4***	0.310***	0.268***	0.400***	0.364***	-0.0371	-0.0442	0.000834	-0.0125
	(17.22)	(20.09)	(0.0547)	(0.0694)	(0.0667)	(0.0849)	(0.0372)	(0.0445)	(0.00782)	(0.00996)
Propensity Score		✓		✓		✓		✓		✓
Observations	2118	1880	2118	1880	2118	1880	1977	1754	2060	1831

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel B: Skill Composition, Labor Share, and Productivity.

	(1)		(2)		(3)		(4)		(5)	
	Productivity		Labor Share		Educ. Years		College Share		Production Worker Share	
Treatment	-0.00742	-0.0148	0.000989	0.00210	-0.0338	-0.0610	-0.00531	-0.00562	0.00735	-0.00679
	(0.0345)	(0.0417)	(0.00500)	(0.00620)	(0.0513)	(0.0649)	(0.00836)	(0.0106)	(0.0181)	(0.0213)
Propensity Score		✓		✓		✓		✓		✓
Observations	2056	1828	2060	1831	1953	1733	1912	1697	1896	1708

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The sample is the main analysis sample (subsidies design) without the balanced-panel requirement. For the firms that exited, the first three outcomes in Panel A are defined as zero, all others are defined as missing.

Figure A-14: Predictive Features for Winning a Technology Subsidy.

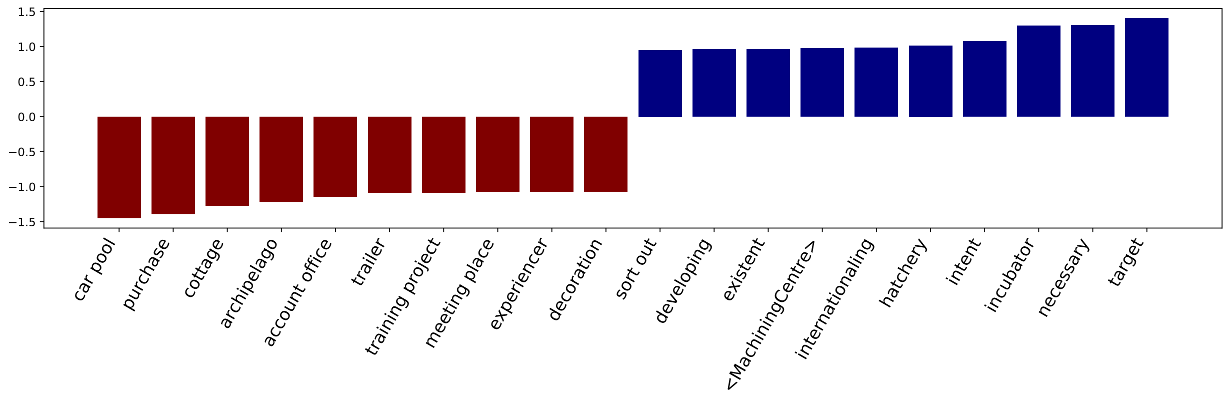

Notes: The features (words) are plotted from top and bottom SVM coefficients predicting treatment status. The y-axis refers to the coefficient size and indicates the relative importance of each feature. Positive (negative) values indicate that the word is typically (not) associated with applications winning a subsidy. The sample is the main analysis sample (subsidies design).

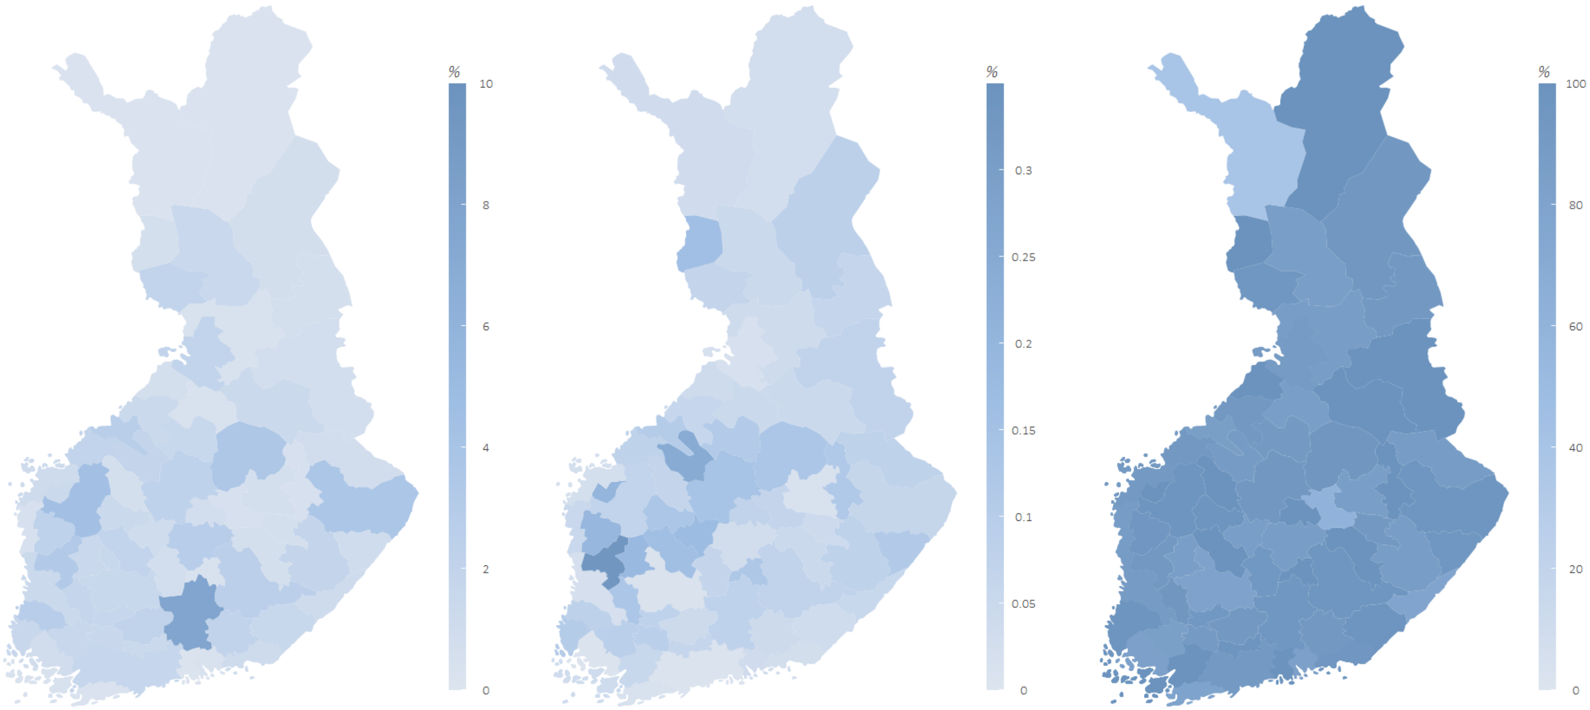

(a) Share of Applications.

(b) Applications per Capita.

(c) Acceptance Rate.

Figure A-15: Maps of Application Statistics.

Notes: The maps visualize descriptive statistics of sample applications (firms) at the subregion level: (a) distribution of the sample over subregions, (b) sample firm count per capita over subregions, i.e., adjusting for the population in each subregion, (c) the win rate by subregion. The sample is the main analysis sample (subsidies design). Applications are more represented in areas with a large manufacturing industry (e.g., the Lahti region). The acceptance rates do not vary significantly by region, barring some outliers with only few applications.

Table A.15: Predicting Treatment.

	(1) Treatment	(2) Treatment	(3) Treatment
Employment (Log)	0.0287***	0.0340***	0.0248***
	(0.00613)	(0.00738)	(0.00706)
Revenue per Worker (EUR 100K)	-0.00850	-0.00924	-0.00462
	(0.00496)	(0.00517)	(0.00456)
Average Wage (EUR 100K)	-0.0251	-0.120	-0.0958
	(0.111)	(0.124)	(0.118)
Profit per Worker (EUR 100K)	-0.108*	-0.127*	-0.0524
	(0.0508)	(0.0546)	(0.0491)
Value Added per Worker (EUR 100K)	0.129*	0.145*	0.0570
	(0.0636)	(0.0686)	(0.0625)
Propensity Score			0.494***
			(0.0645)
Observations	2031	2031	1812

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Coefficients for different variables measured at $\tau = -3$ from OLS estimation, where treatment status is the dependent variable. The sample is the main analysis sample (subsidies design). The propensity score is based on the application texts. The first column reports the coefficients without any controls, the second with a basic set of controls (year indicators and 2-digit industry), and the third when adding text-based propensity score to the controls. Log employment is most predictive of receiving a subsidy: the coefficient is between 0.025-0.034 and is significant in all specifications. Still, the effect is very small, as holding all else equal, doubling the employment of a firm would result in only a 2.0-2.4 percentage points increase in its probability of receiving a subsidy. Interestingly, the first two columns imply that increasing the profit per worker by 10K euros would reduce the probability of a successful application by about 1.1-1.3 percentage points, but a similar increase in value added per worker would increase the probability by about 1.3-1.5 percentage points. Based on these numbers, it seems that high profits themselves would not increase the chances of a successful subsidy application, but a high conversion rate of money spent on materials to value-added increases the probability. Adding the propensity score in Column 3 as a control eradicates the significance of both of these coefficients, leaving only log employment as statistically significant at any of the conventional levels. This implies that the propensity score captures firm characteristics correlated with both profits and value added. Since the coefficient of log employment is so small, and we control for baseline employment in all of our estimates, together with the propensity score, year indicators, and industry controls, the remaining selection bias in our specification and context is potentially minor. The propensity score itself has a highly significant and large coefficient, about 0.49. Thus the score is not perfect in predicting treatment (as it would be if the coefficient were to be unity) but performs reasonably well.

Table A.16: Controlling for the Propensity Score from the Register Data.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Machine Inv.	Employment	Revenue	Wages	Profit	Productivity	Labor Share	Educ. Years	College Share	Prod. Worker Share
Treatment	54.72^{**}	0.310^{***}	0.386^{***}	-0.000362	-0.00297	-0.00571	0.000597	0.0288	0.00657	0.000871
	(17.44)	(0.0594)	(0.0777)	(0.0347)	(0.00788)	(0.0353)	(0.00498)	(0.0616)	(0.00936)	(0.0182)
Propensity Score	2224.7^{***}	-3.276^{***}	-3.017^{***}	-2.127^{***}	0.175^{**}	0.0229	-0.109^{***}	-0.267	-0.0643	-0.00860
	(181.1)	(0.469)	(0.533)	(0.282)	(0.0575)	(0.200)	(0.0320)	(0.524)	(0.0745)	(0.151)
Observations	2031	2031	2031	1952	2031	2031	2031	1884	1884	1891

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: First-difference estimation results on selected outcomes with balance sheet-based propensity score controls. The sample is the main analysis sample (subsidies design). The propensity scores are constructed by first estimating a logit model based on average wages, employment, revenue at $\tau = -3$, and employment and revenue trends from $\tau = -3$ to $\tau = 1$. The predicted treatment probability is then used as the propensity score. Machine investment is summed over $\tau = 0$ to $\tau = 2$ in EUR K. We find qualitatively similar effects as in our baseline estimation: the first stage on machinery investment is clear, employment and revenue grow significantly, but there is no evidence of skill-bias, with fairly precise zeros in most of the other outcomes. Interestingly, the first stage effect is around half in size of that in our preferred estimates (55K vs. 103K EUR). On the other hand, the employment and revenue effects are somewhat larger (about 6 and 14 percentage points respectively) than without the control.

Table A.17: Controlling for Selection Bias Using Qualitative Evaluations.

	(1)		(2)		(3)		(4)		(5)		(6)	
	Machine Inv. (EUR K)		Employment		Revenue		Productivity		Labor Share		College Share	
No good	103.2***	95.38***	0.234***	0.238**	0.298***	0.315**	-0.0104	-0.0122	-0.000850	0.000771	0.00382	0.00329
	(18.43)	(23.23)	(0.0645)	(0.0790)	(0.0823)	(0.102)	(0.0357)	(0.0429)	(0.00509)	(0.00617)	(0.00988)	(0.0123)
No jobs	106.8***	99.24***	0.225***	0.224**	0.306***	0.321**	-0.00710	-0.0109	-0.00240	-0.00110	0.00592	0.00638
	(17.94)	(22.54)	(0.0628)	(0.0768)	(0.0799)	(0.0990)	(0.0353)	(0.0432)	(0.00507)	(0.00620)	(0.00966)	(0.0121)
Propensity Score		✓		✓		✓		✓		✓		✓
N, No-good	2021	1803	2021	1803	2021	1803	2021	1803	2021	1803	1875	1668
N, No jobs	2026	1807	2026	1807	2026	1807	2026	1807	2026	1807	1879	1671

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The effects on selected outcomes when dropping losing (control) firms deemed to not satisfy basic financial requirements ("no-good," n = 10) and not produce enough jobs ("no jobs," n = 5) by the administrative evaluator. The sample is a subset of the main analysis sample (subsidies design). The table addresses the concern of bad counterfactuals explaining our results. As the set of control firms (i.e. those not receiving a subsidy despite applying for one) in our baseline sample is relatively small (n = 146), we are able to read through the evaluation texts written by the program officers for each firm and determine potentially bad counterfactuals. "No-good" refers to the applications where the rejection is due to the firm's poor financial health or other factors indicating that the firm is likely to sustain its business in the long term. "No-jobs" refers to the cases where the officer rejected the application due to it not creating new jobs, a condition not required for acceptance, but in some cases detrimental to the decision. Dropping these potentially problematic control firms does not change the results in any meaningful way.

Table A.18: Propensity-Score Trimmed Samples.

Panel A: Investment, Employment, Wages, and Firm Performance.

	(1) Machine Inv. (EUR K)		(2) Employment		(3) Revenue		(4) Wages		(5) Profit Margin	
5%	112.3***	107.5***	0.245**	0.248**	0.301**	0.327**	-0.0124	-0.00497	-0.00420	-0.00628
	(24.97)	(25.36)	(0.0815)	(0.0828)	(0.116)	(0.118)	(0.0439)	(0.0441)	(0.0118)	(0.0119)
10%	127.4***	123.7***	0.251**	0.254**	0.313*	0.324*	-0.00737	-0.00578	-0.00304	-0.00472
	(28.57)	(28.65)	(0.0953)	(0.0956)	(0.132)	(0.133)	(0.0472)	(0.0473)	(0.0134)	(0.0134)
20%	91.05*	91.01*	0.188	0.188	0.242	0.242	0.00738	0.00652	0.000227	0.000249
	(37.71)	(37.74)	(0.124)	(0.125)	(0.174)	(0.174)	(0.0561)	(0.0559)	(0.0151)	(0.0152)
Propensity Score		✓		✓		✓		✓		✓
N, 5%	1631	1631	1631	1631	1631	1631	1570	1570	1631	1631
N, 10%	1449	1449	1449	1449	1449	1449	1395	1395	1449	1449
N, 20%	1088	1088	1088	1088	1088	1088	1049	1049	1088	1088

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel B: Skill Composition, Productivity, and The Labor Share.

	(1) Productivity		(2) Labor Share		(3) Educ. Years		(4) College Share		(5) Production Worker Share	
5%	-0.0311	-0.0203	0.00105	0.000261	-0.0636	-0.0499	0.00144	0.00226	-0.0316	-0.0309
	(0.0489)	(0.0494)	(0.00677)	(0.00686)	(0.0900)	(0.0911)	(0.0136)	(0.0138)	(0.0227)	(0.0229)
10%	-0.0314	-0.0294	0.000653	0.000813	-0.0273	-0.0251	0.00519	0.00449	-0.0496*	-0.0489
	(0.0553)	(0.0554)	(0.00759)	(0.00765)	(0.102)	(0.103)	(0.0154)	(0.0154)	(0.0252)	(0.0253)
20%	-0.0193	-0.0193	0.00113	0.00106	-0.0281	-0.0286	0.00570	0.00562	-0.0128	-0.0125
	(0.0681)	(0.0682)	(0.00924)	(0.00927)	(0.128)	(0.128)	(0.0180)	(0.0181)	(0.0296)	(0.0296)
Propensity Score		✓		✓		✓		✓		✓
N, 5%	1631	1631	1631	1631	1519	1519	1519	1519	1533	1533
N, 10%	1449	1449	1449	1449	1352	1352	1352	1352	1366	1366
N, 20%	1088	1088	1088	1088	1018	1018	1018	1018	1030	1030

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The estimated effects for the main sample with top and bottom 5%, 10%, and 20% of propensity score values dropped. The results are robust to excluding firms with small and large values of the propensity score. The sample is a subset of the main analysis sample (subsidies design).

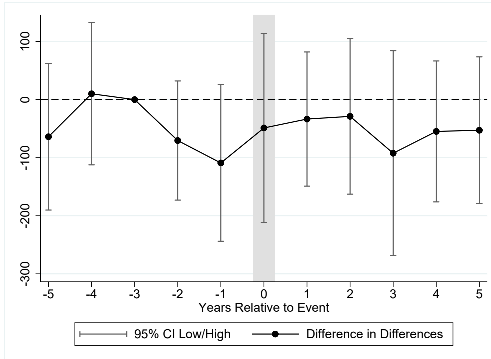

(a) Machinery Investment.

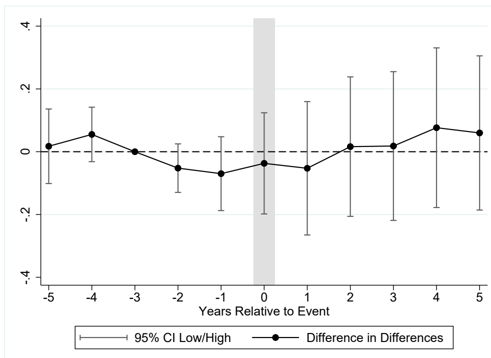

(b) Employment.

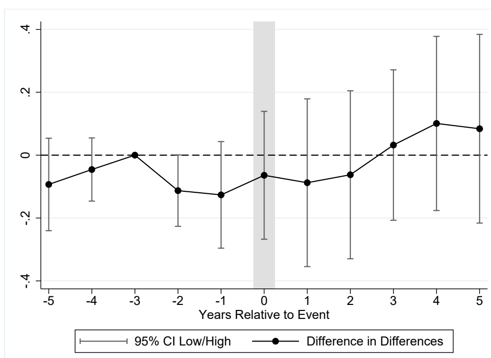

(c) Revenue.

Figure A-16: Placebo Test: The Effects of Insignificant Subsidies.

Notes: Event-study estimates for the effects of insignificant (to the firm) subsidies. The sample includes firms that applied for a subsidy that was less than 10% of their capital stock three years before application. This "placebo test" investigates whether these small subsidies also create treatment effects on machinery investment, employment, and revenue. One concern would that the observed effects in the main sample are coming from the facts winning firms are positively selected: e.g., they are likely to perform better in the future and thus would grow even without the subsidy. If this were true, we would be likely to find positive effects on employment and revenue also in firms where the subsidy itself plays a small role. Reassuringly, we find zero effects on both for all post-application years.

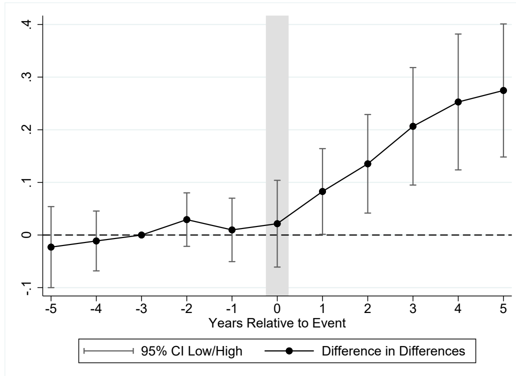

(a) Log Employment.

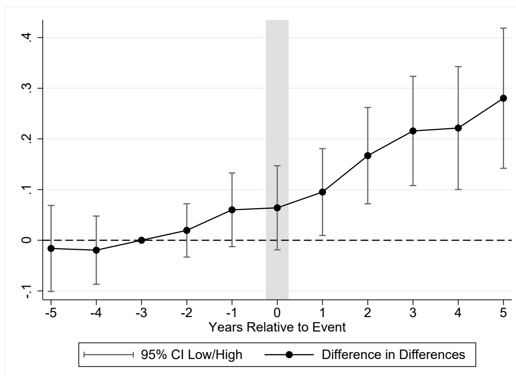

(b) Log Revenue.

Figure A-17: Log Effects.

Notes: The sample is the main analysis sample (subsidies design). Event study graphs of log employment and revenue. The results are very similar to the baseline versions in relative units.

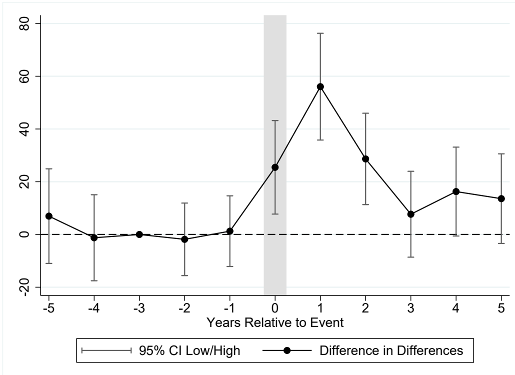

(a) Machinery Investment.

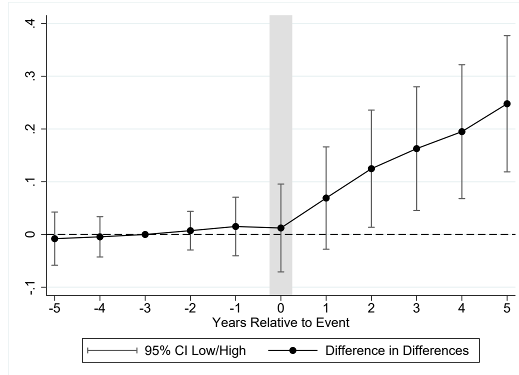

(b) Employment.

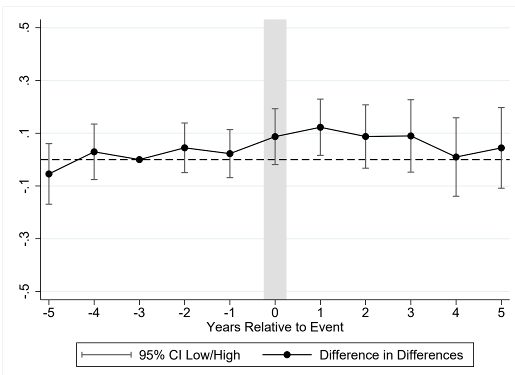

(c) Education Years.

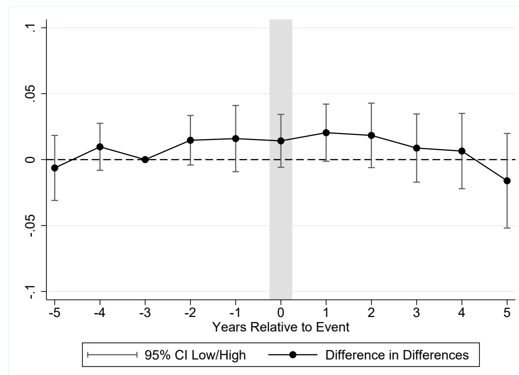

(d) College-Educated Workers' Share.

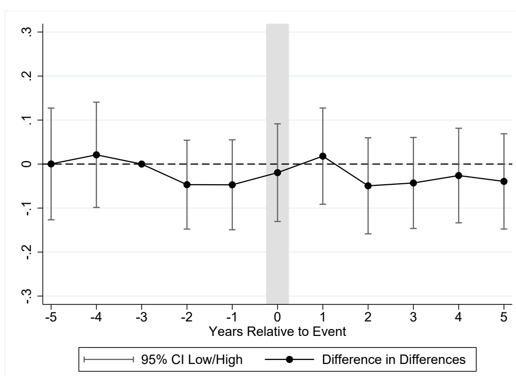

(e) Production Workers' Share.

Figure A-18: Without Controls: Event-Study Estimates of Winners vs. Losers.

Notes: The sample is the main analysis sample (subsidies design). Event study estimates for machinery investment, employment (% relative to $\tau = -3$), average years of education, the employment share of college-educated workers, and the employment share of production workers. No additional controls.

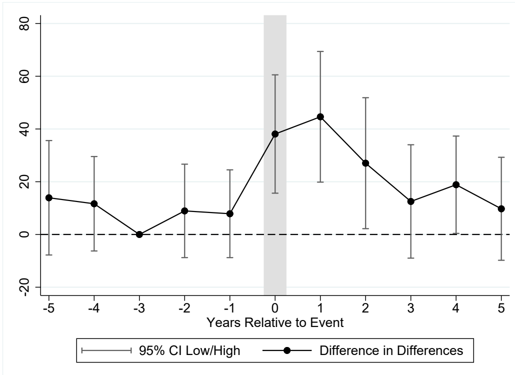

(a) Machinery Investment.

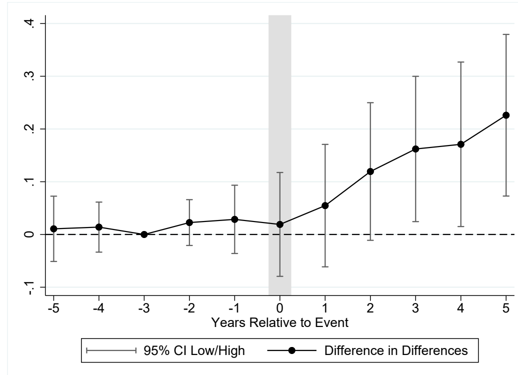

(b) Employment.

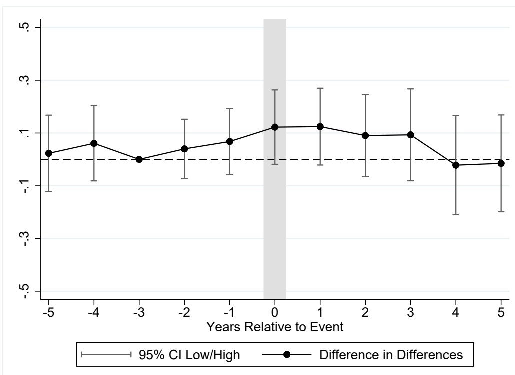

(c) Education Years.

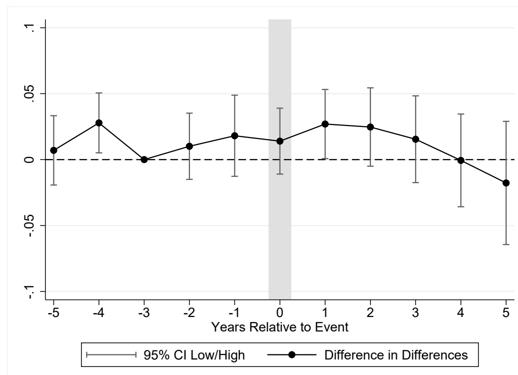

(d) College-Educated Workers' Share.

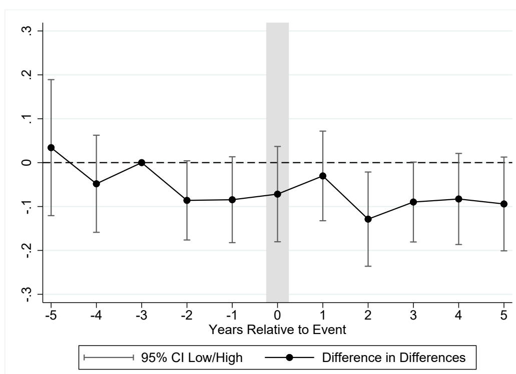

(e) Production Workers' Share.

Figure A-19: With Controls: Event-Study Estimates of Winners vs. Losers.

Notes: The sample is the main analysis sample (subsidies design). Event study estimates for machinery investment, employment (% relative to $\tau = -3$), average years of education, the employment share of college-educated workers, and the employment share of production workers. The specification controls for the firm's employment at $\tau = -3$ interacted with the event-time indicators, 2-digit industry interacted with the calendar-time indicators, and for the text-based propensity score interacted with event-time indicators.

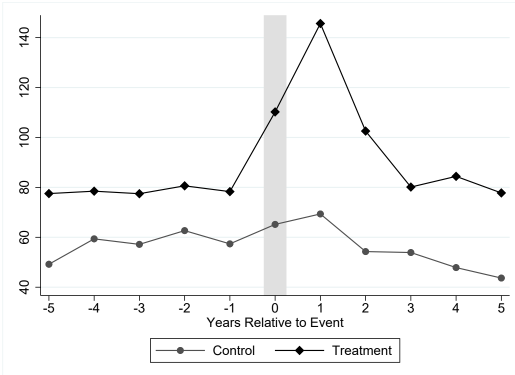

(a) Machinery Investment.

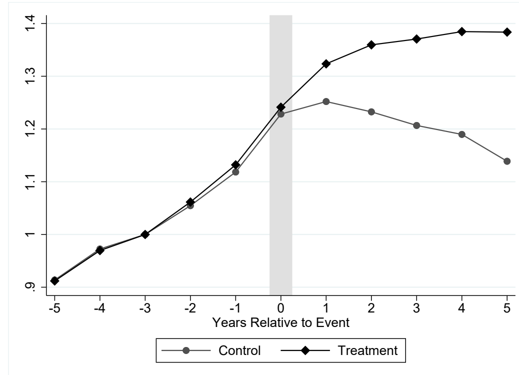

(b) Employment.

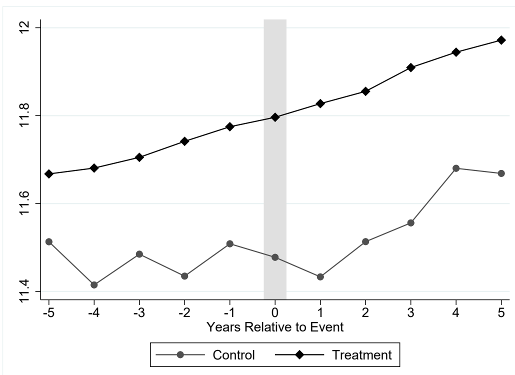

(c) Education Years.

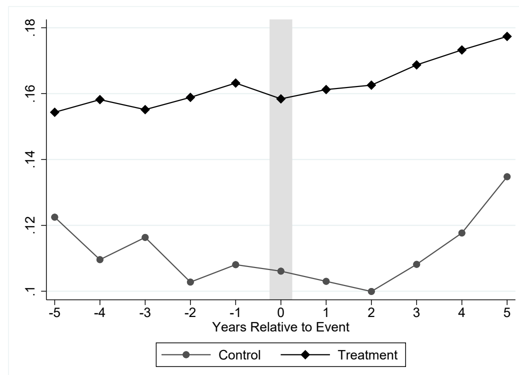

(d) College-Educated Workers' Share.

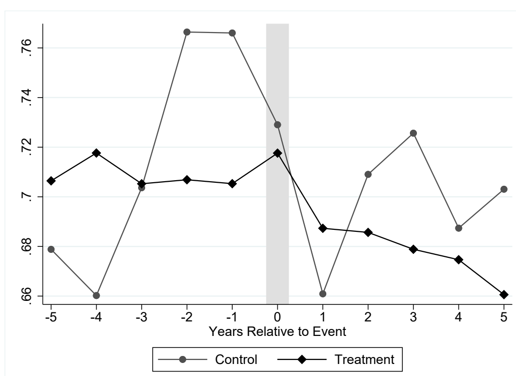

(e) Production Workers' Share

Figure A-20: Raw Means: Winners vs. Losers

Notes: The sample is the main analysis sample (subsidies design). Mean graphs of machinery investment, employment (relative to $t = -3$ level), years in education, the employment share of college-educated workers, and employment share of production workers. Production workers' share is calculated only for firms with more than two full-time employees in a given year.

Table A.19: The Subsidy Program's Maximum Allowed Rates by Area.

Panel A: Years 2000-2006.

Region	Small firm	Medium-sized firm	Large firm
I	40	40	30
II	34	34	25
III	25	25	20
IV	15	0	0

Panel B: Years 2007-2003.

Region	Small firms	Medium-sized firms	Large firms
I	35	25	15
II	25	15	10
III	15	7.5	0
IV	0	0	0

Panel C: Years 2014-2020.

Region	Small firms	Medium-sized firms	Large firms
I	35	25	15
II	30	20	10
III	20	10	0
IV	0	0	0

Notes: The numbers refer to the maximum subsidy rate (%) by area and year. The maximum subsidy rate means the recommended maximum share of the project that the ELY center can subsidize. Source: Finnish Law, sections 1200/200, 1/2007, 675/2007. Accessible at finlex.fi.

Table A.20: Firm Size Definitions.

Size	Workers (M EUR)		Turnover (M EUR)		Balance sheet total (M EUR)
Micro-firm	< 10	and either	≤ 2	or	≤ 2
Small firm (-2007)	< 50	and either	≤ 7	or	≤ 5
Small firm (2007-)	< 50	and either	≤ 10	or	≤ 10
Medium-sized firm	< 250	and either	≤ 50	or	≤ 43
Large firm	≥ 250	or both	≥ 50	and	≥ 43

Notes: The firm size definitions by the EU, used in the subsidy program rules.

Table A.21: Investment and Subsidy Statistics.

Sample Firm Subsidies Share of Manufacturing Investment (All Years)	0.5%
Sample Firm Investment Share of Manufacturing Investment (All Years)	6.9%
Sample Firm Subsidies Share of Manufacturing Investment (Panel Years)	0.4%
Sample Firm Investment Share of Manufacturing Investment (Panel Years)	3.2%
Sample Firm Subsidies Share of Manufacturing Investment (t = 0-2)	0.3%
Sample Firm Investment Share of Manufacturing Investment (t = 0-2)	1.6%
Sample Investment in Total	2,872 M EUR
Manufacturing Investment in Total	93,171 M EUR
Sample Subsidies in Total	320 M EUR
Program Technology Subsidies in Total	758 M EUR
Program Subsidies in Total	2,015 M EUR

Notes: The sample is the main analysis sample (subsidies design). Sample firms are measured at $\tau = -3$, all manufacturing firms cover all possible firm-year combinations that satisfy similar restrictions as the sample firms. The first two numbers represent shares of sample firms appearing in any year, the next two over panel years only ($\tau = -5$ to $\tau = 5$) and the next three over the application year and the two following years ($\tau = 0$ to $\tau = 2$).

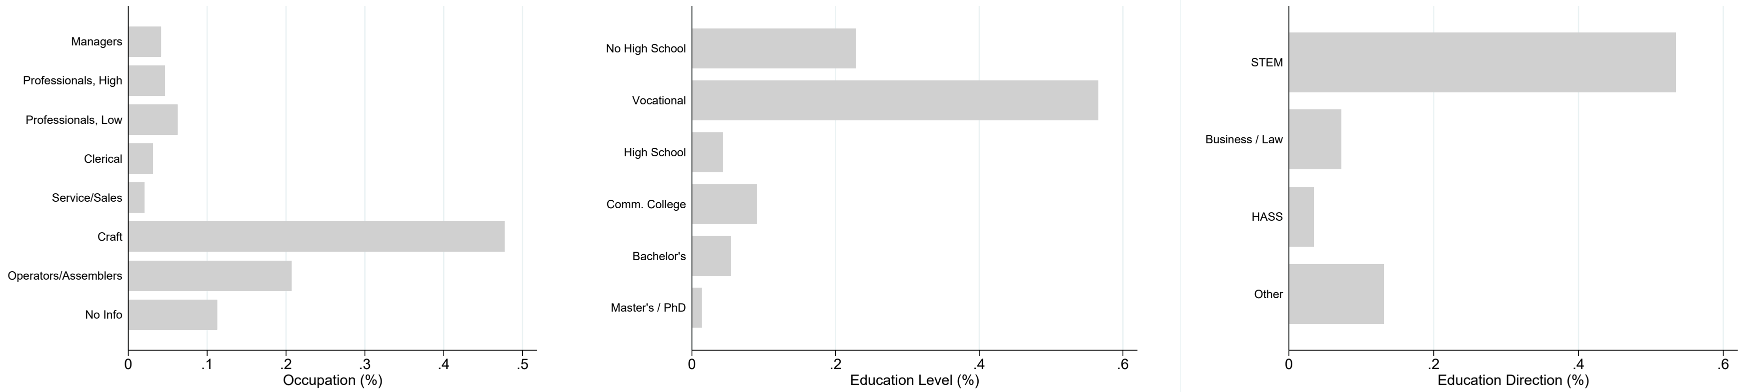

(a) Occupations.

(b) Education Levels.

(c) Education Types.

Figure A-21: Sample Worker's Occupations and Education.

Notes: The figures show the distribution of sample workers' 1-digit occupations, and education levels and types. The sample is the main analysis sample (subsidies design). The shares are unweighted means of the sample firms at $\tau = -3$. We study production work: a vast majority of the workers in the sample firms are craftworkers, operators, and assemblers. The mean share of production workers in the sample firms is approximately 70% of all workers. Notably, the share of clerks and other operation support workers and workers in sales is low. Most workers in the sample hold a vocational school degree or only a primary school degree. The share of workers with a bachelor's degree or higher is low, accounting for less than 20%. A majority, over 50%, of the degrees the workers in the sample firms hold are in STEM fields. Note that the shares in each subfigure do not add up to hundred percent because not all workers have data on occupation or education.

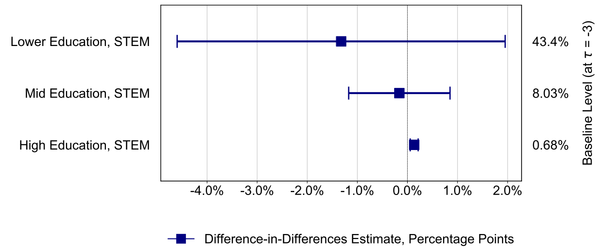

(a) STEM.

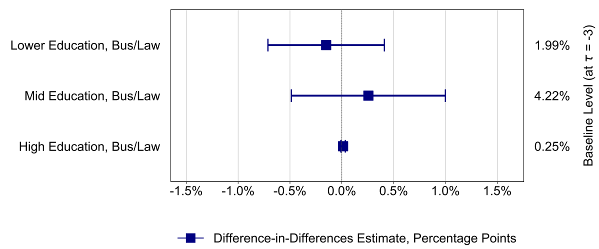

(b) Business and Law.

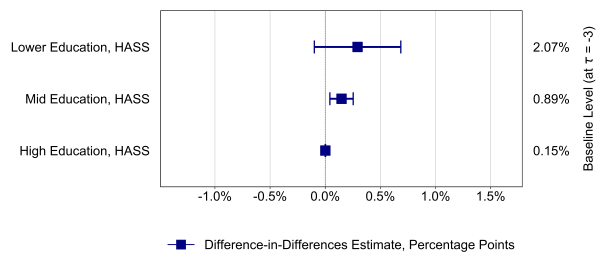

(c) HASS.

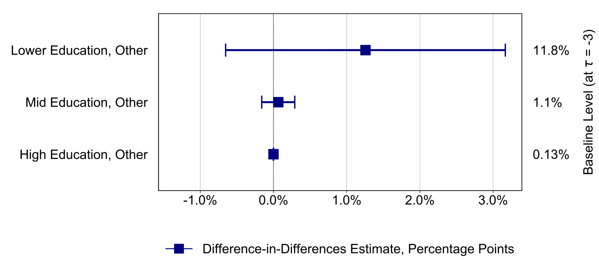

(d) Other.

Figure A-22: Skill Effects: Education by Level and Type.

Notes: The sample is the main analysis sample (subsidies design). The types of education are grouped into science, technology, engineering, mathematics (STEM); business and law; humanities, arts, social sciences (HASS); and others. The levels of education are grouped into lower (high-school or equivalent), mid (BA or equivalent), and high (MA or PhD). We find no economically significant skill composition effects in any of the subgroups of workers. The winning firms increase the share of STEM-educated workers with a Master's or PhD by about 0.15 percentage points. While the effect is very small in the absolute sense, it is significant and translates to around a 20% increase in the group's employment share. There is also a similar effect, about 0.17 percentage points, in the share of HASS-educated workers with a mid-level degree.

Table A.22: Wage Effects by Occupation and Education.

	Occupation			**Education**			
	Prod. Workers	Non-Prod. Low	Non-Prod. High	No Education	Low Education	Mid Education	High Education
Treatment	-0.0342	-0.0746	-0.00812	-0.00776	-0.00774	-0.0428	0.259^{*}
	(0.0237)	(0.0812)	(0.0491)	(0.0470)	(0.0263)	(0.0677)	(0.130)
Baseline	27423.7	25030.4	39791.4	23829.5	24868.1	32054.7	47541.5
N	1833	883	1233	1455	1797	1217	236

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The estimated effects on average wages for different occupation and education groups. The group definitions are given in the main text. The sample is the main analysis sample (subsidies design). The second to last row shows the average levels of wages three years prior to the application for sample firms. The effects are zero for all subgroups of workers other than highly-educated workers (those with MA or PhD degrees). As only 236 firms employ at least one highly-educated worker, the 25.9% positive effect is not necessarily representative of the whole sample. Nonetheless, it could hint toward skill-bias in a subset of the sample firms or rent sharing—part of the increased profits being directed to the owners and executives of the firm, who often are highly educated. In many of the smaller firms, the owners are also employees of the firm.

Table A.23: Wage Effects: Different Wage Outcomes.

Panel A: Level (EUR K).

	(1) Wages (SF; EUR K)		(2) Wages (EUR K)		(3) Wages (Excl. Highest; EUR K)		(4) Highest Wage (EUR K)	
Treatment	0.110	0.340	-0.358	0.217	-0.117	0.386	3.015^{**}	4.250^{***}
	(0.495)	(0.590)	(0.451)	(0.527)	(0.461)	(0.555)	(0.931)	(1.091)
Propensity Score		✓		✓		✓		✓
Baseline	21.95	22.55	25.36	25.61	23.59	23.84	44.92	45.73
N	2031	1812	1884	1676	1766	1577	1884	1676

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel B: Relative (%).

	(1) Wages (SF; %)		(2) Wages (%)		(3) Wages (Excl. Highest; %)		(4) Highest Wage (%)	
Treatment	-0.0481	-0.0285	-0.0293	-0.0000285	-0.0238	0.00106	0.0578^{*}	0.0927^{**}
	(0.0355)	(0.0407)	(0.0239)	(0.0278)	(0.0280)	(0.0341)	(0.0288)	(0.0332)
Propensity Score		✓		✓		✓		✓
Baseline	21.95	22.55	25.36	25.61	23.59	23.84	44.92	45.73
N	1952	1738	1884	1676	1766	1577	1884	1676

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The estimated effects on wage outcomes, both in levels (Panel A) and relative % compared to $\tau = -3$ (Panel B). The baseline means are measured at $\tau = -3$. Treatment is the win-lose indicator. The sample is the main analysis sample (subsidies design). Column 1 wage measure is computed from the firm-level records of Statistics Finland (SF), and the other wage measures are from the worker-level records. The discrepancy between the two wage measures in Columns 1 and 2 are due to the first being the average wage per full-time employee, and the latter per employee headcount. The wage effect is zero in general, but top earner wages appear to grow by about 3 to 4.3 K EUR (or 5.8 to 9.3 percent). Similar to the positive wage effects of highly-educated workers (see Table A.22), one potential reason is that the monetary benefits of the subsidy are directed partly to the wages of the owners and top executives of the firm.

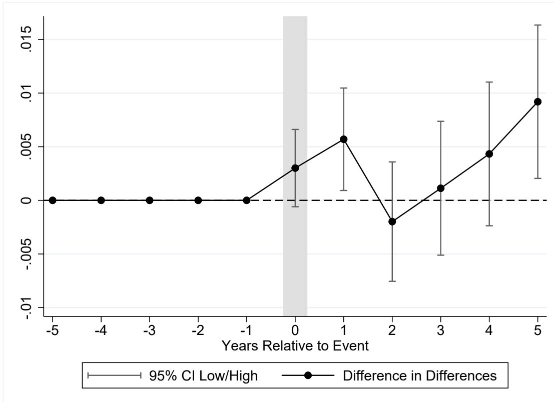

(a) Employed.

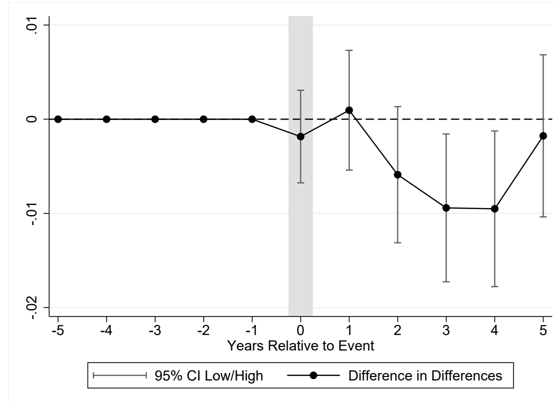

(b) Employed at the Baseline Firm.

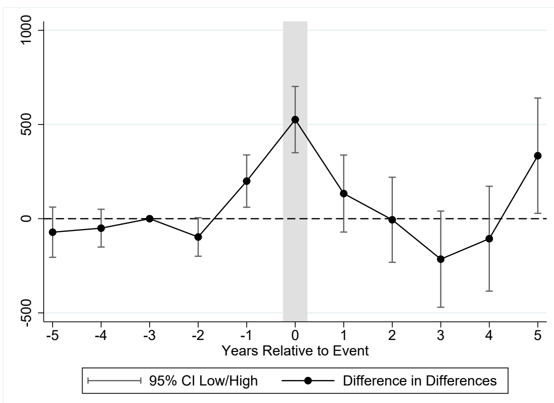

(c) Income.

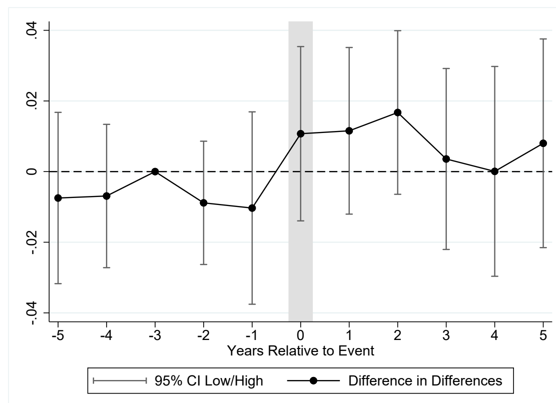

(d) Income Relative to Baseline.

Figure A-23: The Effects on Baseline Workers: Event Studies.

Notes: The sample is the baseline workers (employed at the firm from $\tau = -5$ to $\tau = -1$) in the main analysis sample (subsidies design). The first two and the last outcomes are in percentage points, the third in euros. The baseline workers in treatment group firms are slightly more likely to be employed in general, but less likely to be employed in the baseline firm after the event. The same workers also receive extra income of about 1,000 euros in total in the three years around the application. This corresponds to a salary of about two weeks. One potential reason for this is that the employees could work more hours during the new technology adoption.

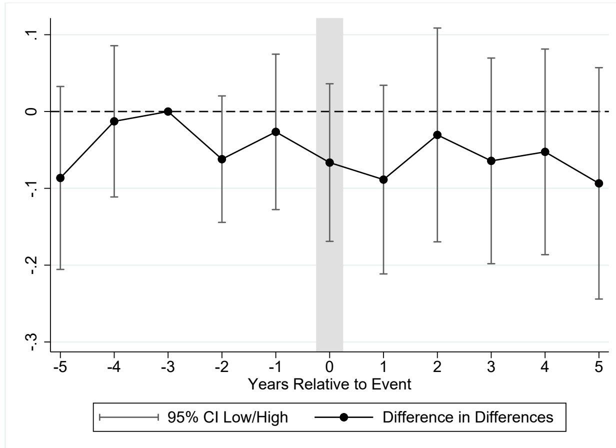

(a) Visuospatial.

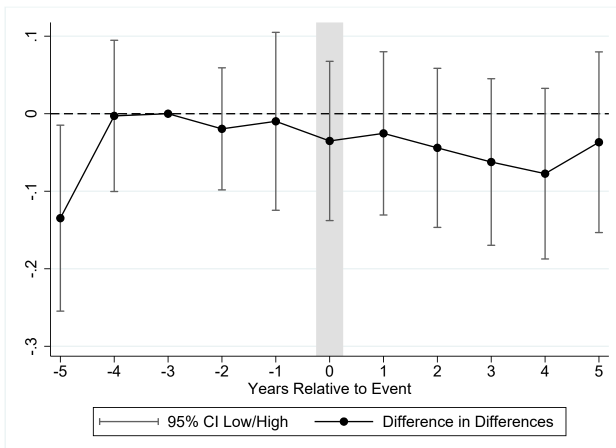

(b) Mathematics.

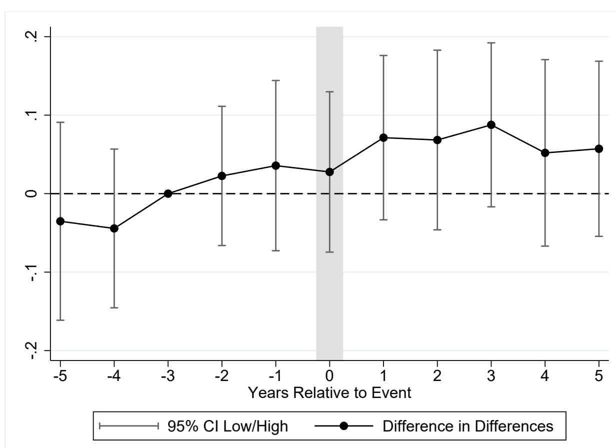

(c) Verbal.

Figure A-24: Skill Effects: Cognitive Performance Event Studies.

Notes: Event study estimates for the average cognitive performance in the firm in standard deviations. The sample is the main analysis sample (subsidies design).

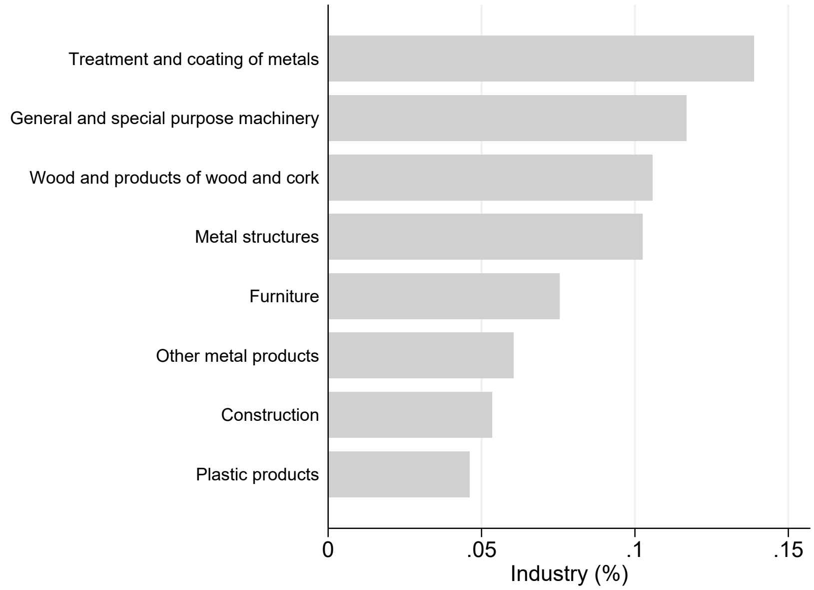

Figure A-25: Sample Firm's Industries.

Notes: The sample firms' top eight industries. The sample is the main analysis sample (subsidies design). The shares are unweighted means of the sample firms at $\tau = -3$. Most firms in the sample operate in metal-related industries, machinery, and construction. Note that the shares do not add up to hundred percent because the industries figure shows only the top eight industries.

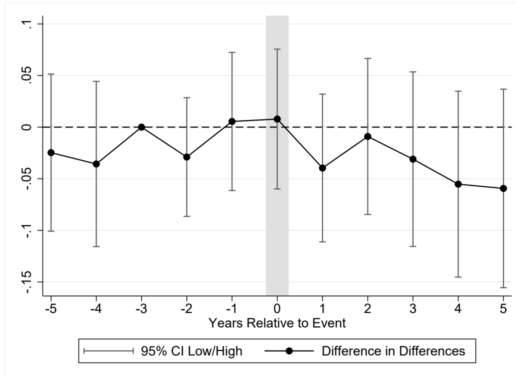

(a) Olley-Pakes.

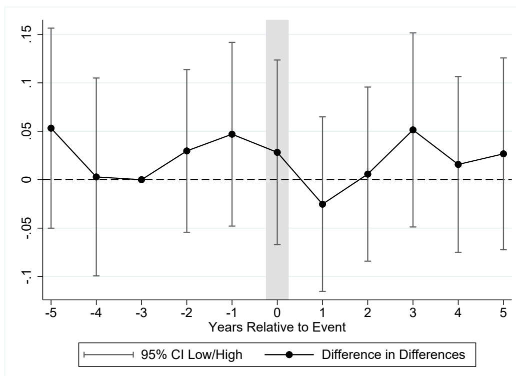

(b) Levinsohn-Petrin.

Figure A-26: Total Factor Productivity: Alternative Versions.

Notes: The sample is the main analysis sample (subsidies design). Event study graphs of log total factor productivity, estimated as in Olley and Pakes (1992) (a) and Levinsohn and Petrin (2003) (b). The results are in line with the Cobb-Douglas version, showing no effect.

Table A.24: The Effects on Capital per Worker.

	(1) Capital per Worker (Win/Lose)	(2) Capital per Worker (Continuous)
Treatment	-84.81	0.00216
	(86.83)	(0.0883)
Baseline	23.85	23845.9
N	1550	1550

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The estimates on capital per worker. Treatment is win/lose status in Column 1 and the amount of subsidies the firm was granted in Column 2. The sample is the main analysis sample (subsidies design). We find no effects.

Table A.25: The Estimated Returns to Capital.

2SLS, Capital Stock Instrumented with Subsidies.

	(1) Gross Profits	(2) Net Profits	(3) Financial Costs
Capital Stock	0.637*	0.256	0.380***
	(0.297)	(0.285)	(0.0539)
Baseline	274006.2	-16074.0	290080.1
N	1560	1560	1560

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The sample is the main analysis sample (subsidies design). 2SLS estimated effects on profits and financial costs. The instrument, dependent value, and outcomes are in EUR. Each column specifies the outcome. Gross and net profits refer to profit before and after financial costs are deducted. The capital stock in euros is instrumented with the amount of subsidies in euros. In principle, the coefficients are interpreted as the response in profits or financial costs.

Table A.26: The Effects of Specific Uses of Technologies on Specific Occupations.

	(1) Worker Share	(2) Wages	(3) Educ. Years	(4) Labor Share
Machining-	0.0454***	0.0753	-0.00340	0.0360*
Machinists	(0.0106)	(0.0816)	(0.0179)	(0.0158)
Baseline	0.0227	25751.0	11.78	0.107
N	554	51	51	51
Welding-	0.0135	0.0598	0.0295	0.00971
Welders	(0.00995)	(0.0938)	(0.0214)	(0.00839)
Baseline	0.0683	25831.3	11.48	0.0594
N	300	88	88	88
Painting-	0.00518	0.0114	0.0258	0.00825
Painters	(0.00358)	(0.0750)	(0.0279)	(0.00590)
Baseline	0.0250	21076.8	10.58	0.0411
N	307	65	65	65
Logistics-	-0.000775	-0.0936	0.0413	-0.00759
Logistics (Non-Office)	(0.000488)	(0.144)	(0.0524)	(0.00699)
Baseline	0.0120	25330.9	10.76	0.0191
N	799	58	58	58
Logistics-	0.0000555	0.298	-0.0884	-0.000707
Logistics (Office)	(0.000246)	(0.301)	(0.112)	(0.00189)
Baseline	0.00399	29623.5	11.70	0.00696
N	799	40	40	40

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The effects of specific technologies on specific workers. The technologies refer to the description of the technology in the subsidy application text. The idea is to see whether a specific technology affects a specific set of workers associated with the technology. The occupations are from the worker's occupational records. The technology-to-worker pairs are: (1) machining-machinists, (2) welding-welders, (3) painting-painters, and (4-5) logistics words (e.g. "driving," "hoisting") to logistics occupations (non-office and office). N refers to the sample size (number of firms) where the given outcome is defined. Note that employment shares are defined for all firms with the given technology, but the other three outcomes require at least one worker with the given occupation, hence the smaller sample size. We find positive effects on the employment share and the wage-bill (labor) share of machinists when the firm has applied for a subsidy specifying the intended use to be an investment in technologies associated with machining. These effects are sizable, considering the relatively small baseline share of workers with the occupational title "machinist" employed. The sample is winners matched to non-applicants (the matching procedure described in the paper).

Table A.27: The Effects of Specific Types and Uses of Technologies Measured from the Text Data.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Sh.	(7) Educ. Years	(8) College Sh.	(9) Prod. Sh.	(10) Obs.
Types of Technology										
CNC	158.0***	0.168***	0.180***	0.0367	-0.0100	0.00299	-0.0331	0.00304	-0.0211	628
	(11.76)	(0.0398)	(0.0492)	(0.0212)	(0.0266)	(0.00498)	(0.0513)	(0.00771)	(0.0144)	
Robot	294.4***	0.233*	0.416**	-0.0380	0.0124	-0.0133	0.0311	0.00256	0.0198	232
	(53.60)	(0.0921)	(0.136)	(0.0291)	(0.0593)	(0.00787)	(0.0581)	(0.00961)	(0.0192)	
Laser	164.1***	0.322**	0.313*	-0.0272	-0.0831	0.000777	0.0578	0.0180	0.0139	224
	(38.69)	(0.0979)	(0.132)	(0.0415)	(0.0501)	(0.0101)	(0.0919)	(0.0157)	(0.0286)	
Uses of Technology										
Machining	227.7***	0.246***	0.276***	-0.00449	-0.0187	-0.00126	0.0136	0.00862	-0.0121	584
	(22.60)	(0.0471)	(0.0663)	(0.0227)	(0.0302)	(0.00514)	(0.0505)	(0.00899)	(0.0137)	
Welding	109.1***	0.352***	0.385***	-0.00611	-0.0146	-0.0100	0.0185	-0.00120	0.00966	312
	(18.61)	(0.0835)	(0.0821)	(0.0352)	(0.0431)	(0.00760)	(0.0700)	(0.0122)	(0.0218)	
Painting	161.9***	0.267***	0.318***	-0.0223	0.00608	-0.00591	-0.0112	-0.00147	-0.00148	312
	(28.80)	(0.0634)	(0.0836)	(0.0302)	(0.0396)	(0.00636)	(0.0627)	(0.00946)	(0.0193)	
Logistics	162.2***	0.304***	0.404***	0.00781	0.0348	-0.00659	0.0207	0.0148*	-0.0106	822
	(16.02)	(0.0404)	(0.0544)	(0.0171)	(0.0255)	(0.00388)	(0.0370)	(0.00611)	(0.0108)	
Automation	177.2***	0.178***	0.217***	0.00216	0.0249	-0.00237	0.0546	0.00942	0.000592	678
	(22.54)	(0.0350)	(0.0446)	(0.0189)	(0.0259)	(0.00422)	(0.0391)	(0.00650)	(0.0113)	

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The effects of specific types and uses of technology on selected outcomes for the relevant subsets of treatment firms matched to non-applicant control firms. The first three rows show the effects for treatment firms that intend to buy the given technology, and the latter four are the firms that specify the listed uses for the technologies. The "Obs." column refers to the sample size: the number of firms with subsidy application texts containing keywords associated with the given technology or its use. Machine investment is in EUR K. The broad interpretation of the results is that the firm-level effects do not vary significantly across the specified technologies or their uses. The sample is the winners matched to non-applicants (the matching procedure described in the paper).

Table A.28: Hardware vs. Software Events.

	(1) Machine Inv. (EUR K)		(2) Employment		(3) Revenue		(4) Productivity		(5) Labor Share		(6) College Share		(7) Obs.
Hardware (Register)	81.52***	80.72***	0.216***	0.212**	0.299***	0.308**	-0.000775	-0.00660	-0.00138	-0.0000368	0.00390	0.00548	1,726
	(15.60)	(20.05)	(0.0625)	(0.0754)	(0.0791)	(0.0970)	(0.0354)	(0.0433)	(0.00502)	(0.00607)	(0.00946)	(0.0118)	
Software (Register)	281.8***	260.1***	0.296***	0.330**	0.361***	0.388**	-0.0584	-0.0553	-0.00298	-0.00132	0.0189	0.0109	451
	(30.10)	(38.55)	(0.0826)	(0.105)	(0.102)	(0.134)	(0.0431)	(0.0534)	(0.00585)	(0.00741)	(0.0115)	(0.0148)	
Hardware (Text)	105.7***	99.68***	0.234***	0.238**	0.312***	0.337***	-0.0137	-0.0101	-0.00151	-0.000185	0.00527	0.00599	1,971
	(18.13)	(22.49)	(0.0634)	(0.0761)	(0.0804)	(0.0979)	(0.0360)	(0.0435)	(0.00506)	(0.00610)	(0.00959)	(0.0116)	
Software (Text)	-19.96	-9.464	0.329	0.450	0.573	0.873	-0.113	0.00135	0.00929	-0.0168	-0.0286	-0.0212	107
	(178.3)	(207.5)	(0.253)	(0.371)	(0.392)	(0.584)	(0.237)	(0.277)	(0.0264)	(0.0308)	(0.0457)	(0.0580)	
Propensity Score		✓		✓		✓		✓		✓		✓	

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The effects of hardware and software events. The sample is a subset of the main analysis sample (subsidies design). The first two rows include the treatment firms that have software purchases in the three-year period after the application (software row) and those that did not (hardware row). These groups are mutually exclusive. The latter two rows include firms stating the intention to purchase hardware or software technologies in the application texts. They both include all losing (control group) firms. The text-based categories are not mutually exclusive, so that an application can include both intended hardware and software investments and thus appear in both categories. The results for text-based software events are highly imprecise, largely due to the small sample size (n = 107). The register-based classification implies that subsidies associated with software purchases induce larger investment and lead to larger employment and revenue effects.

Table A.29: The Effects by Industry Type: Automation, Skill-Level, and Tradability.

Panel A: High vs. Low Automation.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Share	(7) Educ. Years	(8) College Share	(9) Prod. Work. Share
High Automation	115.8***	0.189	0.212	-0.0369	-0.0900	0.00405	0.0675	0.000842	-0.00197
	(25.96)	(0.104)	(0.122)	(0.0481)	(0.0516)	(0.00702)	(0.0866)	(0.0149)	(0.0246)
Low Automation	143.0***	0.277*	0.295	-0.0000105	-0.0568	0.00835	0.0687	0.0211	-0.00410
	(31.95)	(0.109)	(0.156)	(0.0695)	(0.0805)	(0.00907)	(0.121)	(0.0166)	(0.0451)
N, High Automation	1223	1223	1223	1179	1223	1223	1136	1136	1142
N, Low Automation	474	474	474	457	474	474	448	448	443

Panel B: High vs. Low Skill.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Share	(7) Educ. Years	(8) College Share	(9) Prod. Work. Share
High Skill	114.4***	0.255	0.445***	-0.121	-0.0881	-0.00237	0.225	0.0526*	0.0438
	(31.79)	(0.130)	(0.100)	(0.0901)	(0.0790)	(0.0103)	(0.127)	(0.0224)	(0.0472)
Low Skill	103.1***	0.225**	0.282**	-0.0263	0.0148	-0.00143	-0.0289	-0.00680	-0.00962
	(20.32)	(0.0699)	(0.0946)	(0.0380)	(0.0393)	(0.00571)	(0.0690)	(0.00985)	(0.0195)
N, High Skill	532	532	532	511	532	532	499	499	497
N, Low Skill	1499	1499	1499	1441	1499	1499	1385	1385	1394

Panel C: Tradable vs. Non-Tradable Industries.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Share	(7) Educ. Years	(8) College Share	(9) Prod. Work. Share
Tradable	130.0***	0.230**	0.298**	-0.0430	-0.0451	0.00132	0.0852	0.00939	0.0146
	(22.86)	(0.0837)	(0.101)	(0.0474)	(0.0436)	(0.00637)	(0.0747)	(0.0128)	(0.0239)
Non-Tradable	70.58**	0.234**	0.334**	-0.0537	0.0525	-0.00632	-0.0480	0.00113	-0.0186
	(26.59)	(0.0902)	(0.123)	(0.0533)	(0.0573)	(0.00791)	(0.103)	(0.0136)	(0.0280)
N, Tradable	1509	1509	1509	1450	1509	1509	1402	1402	1404
N, Non-Tradable	522	522	522	502	522	522	482	482	487

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The sample is a subset of the main analysis sample (subsidies design). Estimated effects on selected outcomes for firms in high vs. low automation (Panel A), high vs. low skill industry (Panel B), and tradable vs. non-tradable output industry (Panel C). The division with respect to automation level is defined by using classifications in Acemoglu and Restrepo (2020) harmonized to the Finnish industries. An industry is classified into high skill if it is above the median industry in average years of education of its workers, and low skill if below. Tradability is defined by using classifications in Mian and Sufi (2014) harmonized to the Finnish industries. Machine investment is in EUR K.

Table A.30: The Effects by Industry Type With Propensity Score Controls: Automation, Skill-Level, and Tradability.

Panel A: High vs. Low Automation.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Share	(7) Educ. Years	(8) College Share	(9) Prod. Work. Share
High Automation	101.1**	0.114	0.213	-0.0217	-0.0337	0.000346	0.0656	0.00420	-0.0292
	(32.45)	(0.124)	(0.146)	(0.0509)	(0.0574)	(0.00827)	(0.110)	(0.0190)	(0.0291)
Low Automation	135.3**	0.385**	0.370*	-0.00960	-0.136	0.0155	-0.0423	0.0146	-0.0253
	(43.81)	(0.129)	(0.186)	(0.0892)	(0.111)	(0.0115)	(0.148)	(0.0186)	(0.0561)
N, High Automation	1098	1098	1098	1055	1098	1098	1016	1016	1026
N, Low Automation	414	414	414	399	414	414	391	391	390

Panel B: High vs. Low Skill.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Share	(7) Educ. Years	(8) College Share	(9) Prod. Work. Share
High Skill	102.5*	0.415*	0.624***	-0.0660	-0.129	0.00306	0.379	0.0802*	0.0252
	(45.34)	(0.180)	(0.177)	(0.113)	(0.106)	(0.0136)	(0.198)	(0.0331)	(0.0657)
Low Skill	95.39***	0.197*	0.278*	-0.0163	0.0175	-0.000739	-0.0796	-0.00951	-0.0277
	(24.74)	(0.0830)	(0.111)	(0.0434)	(0.0464)	(0.00671)	(0.0796)	(0.0116)	(0.0220)
N, High Skill	461	461	461	442	461	461	431	431	433
N, Low Skill	1351	1351	1351	1296	1351	1351	1245	1245	1259

Panel C: Tradable vs. Non-Tradable Industries.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Share	(7) Educ. Years	(8) College Share	(9) Prod. Work. Share
Tradable	121.4***	0.190	0.283*	-0.0409	-0.0476	-0.000244	0.0813	0.00877	-0.0220
	(29.74)	(0.107)	(0.132)	(0.0538)	(0.0559)	(0.00784)	(0.0987)	(0.0168)	(0.0285)
Non-Tradable	62.85	0.293**	0.387**	-0.00810	0.0505	0.000338	-0.102	0.00436	-0.0201
	(32.72)	(0.100)	(0.135)	(0.0653)	(0.0663)	(0.00942)	(0.117)	(0.0160)	(0.0323)
N, Tradable	1344	1344	1344	1289	1344	1344	1243	1243	1254
N, Non-Tradable	468	468	468	449	468	468	433	433	438

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The sample is a subset of the main analysis sample (subsidies design). Estimated effects on selected outcomes for firms in high vs. low automation (Panel A), high vs. low skill industry (Panel B), and tradable vs. non-tradable output industry (Panel C). The division with respect to automation level is defined by using classifications in Acemoglu and Restrepo (2020) harmonized to the Finnish industries. An industry is classified into high skill if it is above the median industry in average years of education of its workers, and low skill if below. Tradability is defined by using classifications in Mian and Sufi (2014) harmonized to the Finnish industries. Machine investment is in EUR K. Propensity score is included as control.

Table A.31: The Effects by Management Scores.

Panel A: Investment, Employment, Wages, and Firm Performance.

	(1) Machine Inv. (EUR K)	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity
High Score	465.5***	0.476***	0.644***	-0.0316	0.0700
	(97.44)	(0.0953)	(0.128)	(0.0231)	(0.0454)
Low Score	298.6	0.518**	0.425*	0.0112	-0.00715
	(188.7)	(0.180)	(0.174)	(0.0363)	(0.0819)
N, High Score	184	184	184	184	184
N, Low Score	80	80	80	80	80

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel B: Skill Composition and The Labor Share.

	(1) Labor Share	(2) Educ. Years	(3) College Share	(4) Production Worker Share
High Score	-0.0113	-0.0215	-0.00426	0.0140
	(0.00723)	(0.0478)	(0.00870)	(0.0157)
Low Score	0.000285	-0.132	-0.00919	-0.0117
	(0.0106)	(0.109)	(0.0168)	(0.0248)
N, High Score	184	184	184	184
N, Low Score	80	80	80	80

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The sample is a subset of the main analysis sample (subsidies design). Estimated effects on selected outcomes for firms with high vs. low management score, measured using the FMOP as surveyed and defined in Ohlsbom and Maliranta (2021).

Figure A-27: Predictive Features for Text Categories: Process and Product

Panel A: Process.

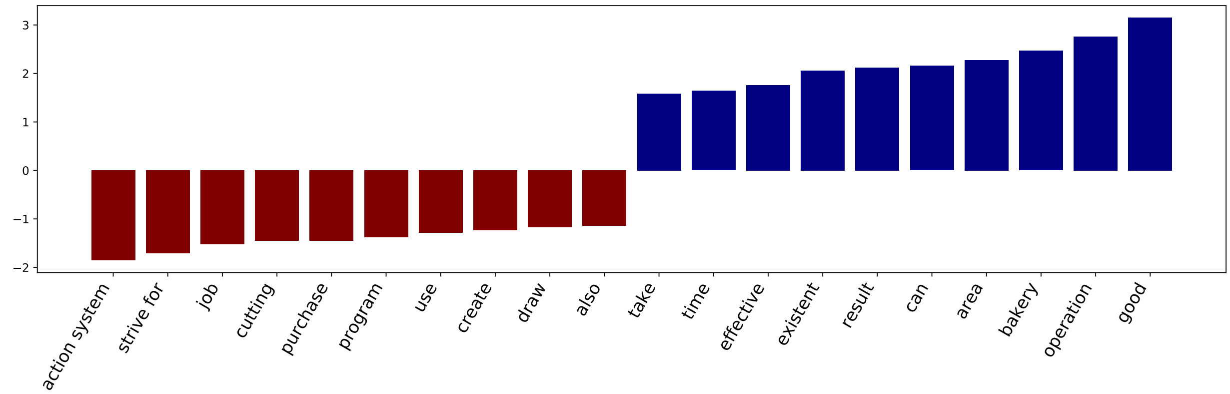

Panel B: Product.

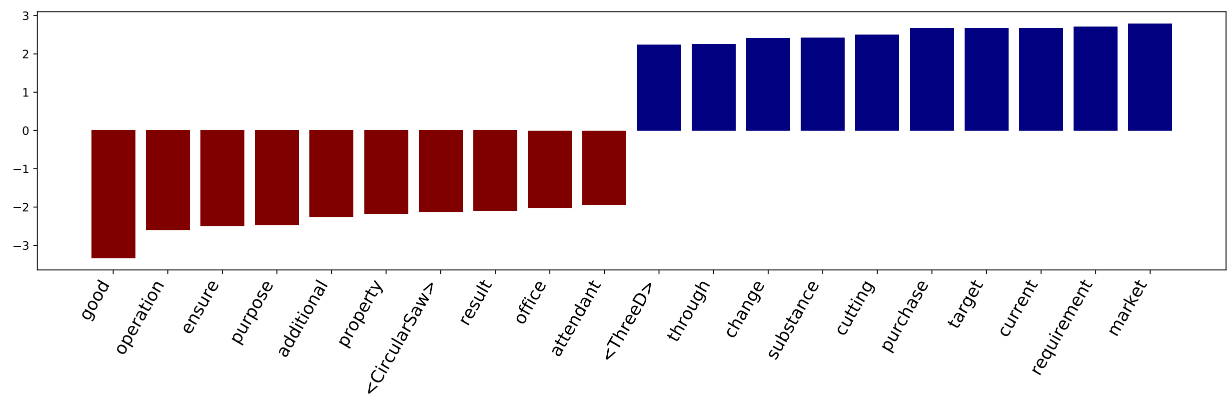

Notes: The features (words) are plotted from top and bottom SVM coefficients predicting the two uses of technologies. The y-axis refers to the coefficient size, and it measures the relative importance of each feature. Positive (negative) values indicate that the word is typically (not) associated with applications in the category.

Table A.32: Continuous Treatment Estimates by Text Categories.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Share	(7) Educ. Years	(8) College Share	(9) Production Share
Product	1.323***	0.194***	4.585***	-1.663	150.7	-0.000345*	0.000338	0.000296	-0.0000846
	(0.0865)	(0.0239)	(0.613)	(9.767)	(124.2)	(0.000130)	(0.00113)	(0.000192)	(0.000323)
Process	1.243***	0.194	1.115	26.97	-359.8	0.000615	-0.00567	-0.00138	0.00123
	(0.216)	(0.109)	(2.399)	(40.53)	(603.8)	(0.000534)	(0.00739)	(0.00144)	(0.00135)
N, Product	2046	2046	2046	2046	2046	2046	1905	1905	1921
N, Process	198	198	198	198	198	198	186	186	186

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The sample is a subset of the main analysis sample (subsidies design). Product type changes refers to technology projects that aim to produce a new type of output. Process type change refers to technology projects that aim to produce the same type of output with the new technologies. Columns 1 and 3 are in EUR. Treatment is scaled to EUR 10,000 for rest of the columns. Columns 6, 8, and 9 (shares) are in percentage points. Column 7 (education years) is in years. Machine investment is the sum over $\tau \in [0, 2]$. Other outcomes are averages over $\tau \in [2, 5]$. N refers to matched observations (matching procedure is described in the paper).

Table A.33: Export Products' and Regions' Skill Intensity.

	(1) Product Skill Intensity	(2) Region Skill Intensity
Treatment	-0.0267	-0.00139
	(0.0599)	(0.0316)
Baseline	12.64	12.87
N	401	401

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The effects on export product and region skill intensity for the sample firms that export. The sample is a subset of the main analysis sample (subsidies design). To construct the outcomes, we first take the average worker education years for each export product and region by taking the average over all years from firms that export the given product and export to the given region. Then for each exporting firm in our sample, we calculate the skill intensity each year by taking the unweighted average of the skill intensities of the products the firm exports that year or the regions it exports to. Export regions and products are measured from the Finnish Customs' Foreign Trade Statistics. A concern about the lack of skill-bias effects in our sample is that it exists, but is subtle and hard to find empirically. One way to explore this possibility is to estimate whether, after adopting new technologies, the firms export products which require more skills or export to regions that do. If this is true, the firms are likely also to exhibit an increased need for skills, even if we do not detect these effects in the short term. This table explores these effects on export products' and regions' skill intensity. The coefficients on both outcomes are fairly precise zeros, implying that the hypothesis of undetected skill bias through this channel does not receive support.

Table A.34: Exporter and Non-Exporter Firms' Skill Intensity.

	Sample; Non-Exporters	Sample; Exporters	Manufacturing; Non-Exporters	Manufacturing; Exporters
Educ. Years	11.66	12.03	11.54	12.30
Firm-Year Observations	1,390	641	218,945	41,275
Firm Observations	1,390	641	16,437	2,102

Notes: Descriptive statistics on the exporter and non-exporter firms' skill intensity. The sample is the main analysis sample (subsidies design) and Finnish manufacturing. The table reports mean worker education years for (1) sample firms that do not export, (2) sample firms that export, (3) all manufacturing firms that do not export, and (4) all manufacturing firms that export. Sample firms are measured at $\tau = -3$, all manufacturing firms cover all possible firm-year combinations that satisfy similar restrictions as the sample firms. Export status is measured using the definition by Statistics Finland. A firm is defined as an exporter in a given year if its total export value is over 12K EUR during the calendar year spread over at least two different months, or a single export event is over 120K EUR in value. Exporting firms in both groups employ more educated workers, confirming a common proposition that exporting firms are more skill intensive. Notably, this difference is smaller in the analysis sample than in manufacturing.

Table A.35: Import and Input Outcomes.

	(1) Import Value (EUR K)	(2) Import Share	(3) Machine Import Value (EUR K)	(4) Machine Import Share	(5) Input Value (EUR K)	(6) Input Share
Treatment	20.60^{***}	0.00287^{*}	3.437^{***}	0.000452	-115.6	-0.0521
	(5.989)	(0.00136)	(1.031)	(0.000249)	(663.6)	(0.0438)
Baseline	152.9	0.0203	27.80	0.00373	3457.9	0.292
N	2031	2031	2031	2031	321	321

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The sample is the main analysis sample (subsidies design). Effects on total value and shares (total value share of revenue) for imports and inputs. Machinery imports include only imports we classified as machinery based on customs codes. The baseline means are measured at $\tau = -3$. The results show that the effects on both machinery imports and all imports (including machinery imports) are positive and significant: machinery imports increase by about 3.4K euros and all imports by about 20.6K euros. While these effects are small, it is important to note that only a small fraction of the sample firms import at all during the panel years. Thus it is less surprising that the average effect is smaller. Similarly, we detect small effects on import share of revenue: about 0.3 percentage points for all imports and zero on machinery import share. The estimated effects on input outcomes (including non-imports) are imprecise, but close to zero. This is partly due to a small sample size (n = 321), as we only observe input outcomes for a subset of firms that have answered the manufacturing survey issued by Statistics Finland.

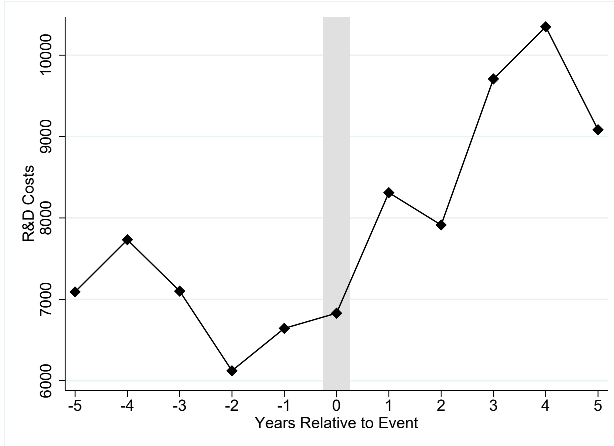

Figure A-28: R&D Expenditure.

Notes: The raw means of R&D expenditure for the subsidy applicant firms, both treatment and control. The sample is the main analysis sample (subsidies design).

A.2 Winners-Losers: Matched Control Group

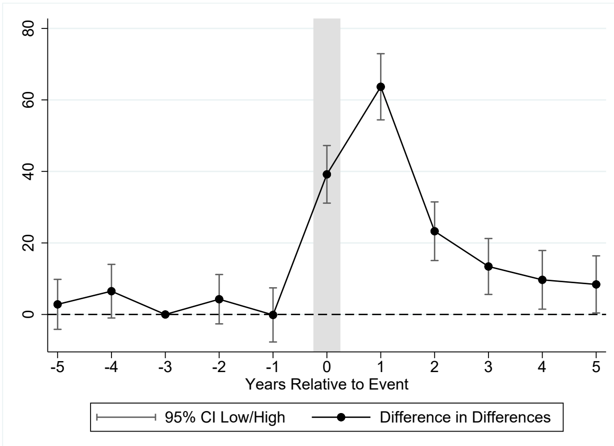

(a) Winners vs. Matched Control:
Machinery Investment.

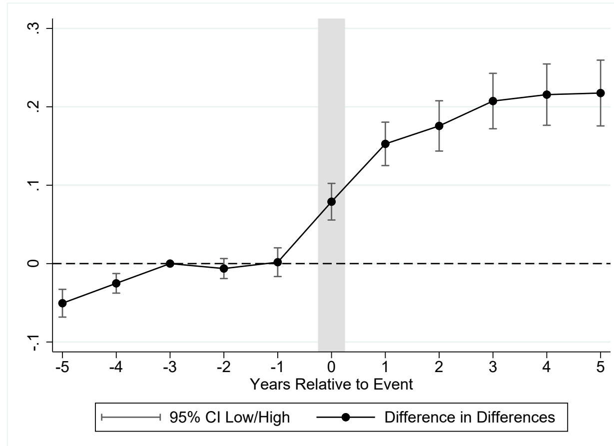

(b) Winners vs. Matched Control:
Employment.

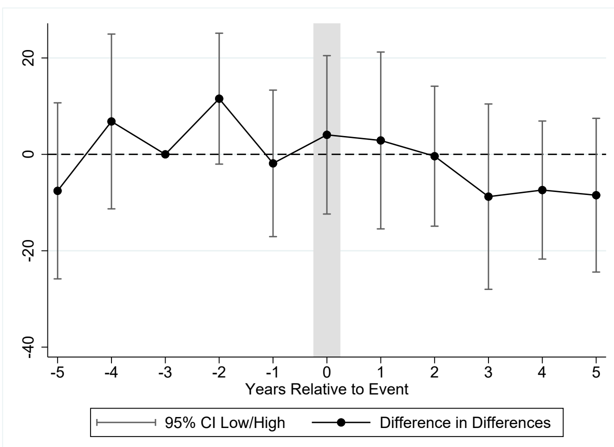

(c) Losers vs. Matched Control:
Machinery Investment

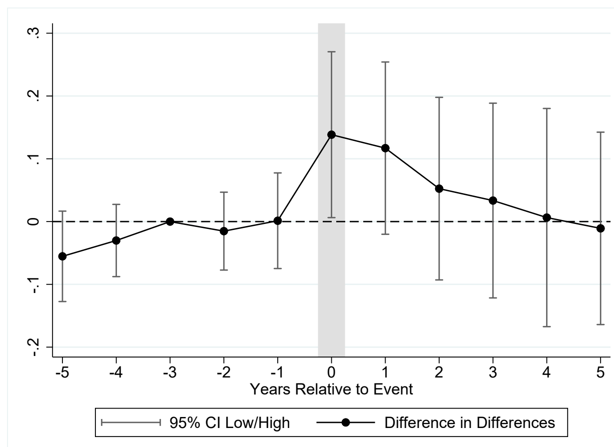

(d) Losers vs. Matched Control:
Employment.

Figure A-29: The Matched Control Groups: The First Stage and Employment Effects.

Notes: Event-study estimates from Equation 1.1. **Panels (a, b):** Treatment group is the subsidy winners (the main treatment group), and control group is constructed via matching. **Panels (c, d):** Treatment group is the subsidy losers (the main control group), and the control group is constructed via matching, i.e., comparing two different control groups. We use coarsened exact matching (CEM). We match by revenue, employment, wages at $\tau = -3$ plus revenue and employment changes in percentages from $\tau = -3$ to $\tau = -1$ and industries' main sectors (letter classes). The CEM percentiles are 10, 25, 50, 75, 90, and 99. The match is 1:1 with replacement. Event time $\tau = 0$ refers to the application year. Back to Section 1.5.

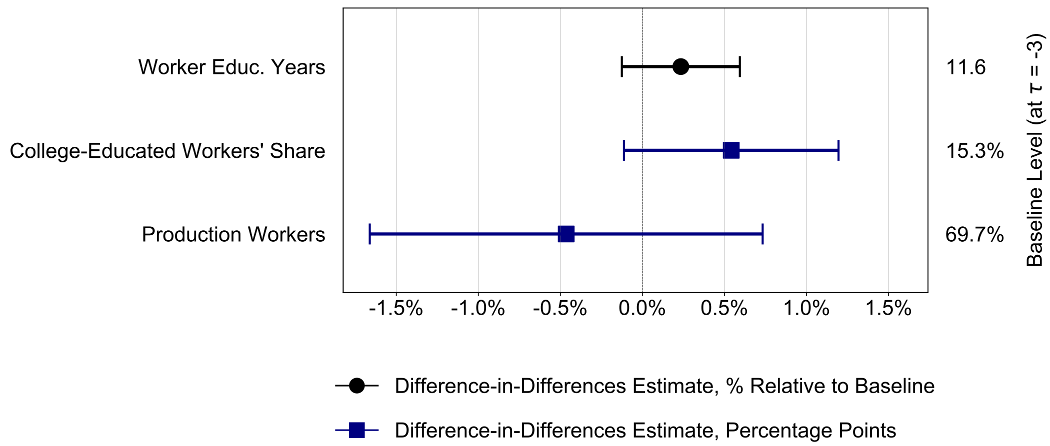

Figure A-30: The Matched Control Group: Skill Effects.

Notes: Difference-in-differences estimates from Equation 1.2. The estimates compare the main treatment group ("winners") to a matched control group. The right-hand side reports means at $\tau = -3$. Back to Section 1.5.

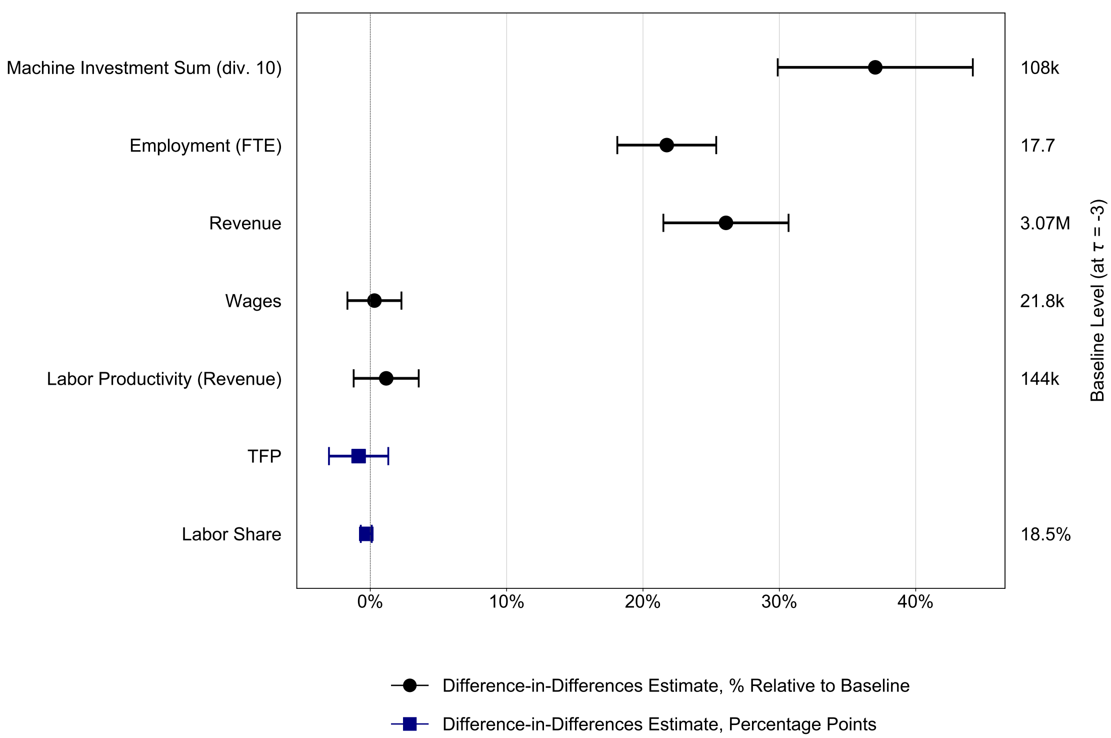

Figure A-31: The Matched Control Group: Firm-Level Effects.

Notes: Difference-in-differences estimates from Equation 1.2. The estimates compare the main treatment group ("winners") to a matched control group. The right-hand side reports means at $\tau = -3$. Back to Section 1.5.

Table A.36: The Matched Control Group: Balance Table A (Winners vs. Matched Control).

	Treatment Group		**Control Group**		**Both**		
Variable	Mean	Std. Dev.	Mean	Std. Dev.	10p	Median	90p
Revenue (EUR M)	3.06	26.57	3.09	9.15	0.17	0.96	6.26
Employment	17.46	46.27	18.03	38.79	1.60	8.20	37.70
Wages (EUR K)	21.60	8.08	22.06	8.36	12.15	22.43	30.56
Subsidy Applied (EUR K)	108.52	126.79	0.00	0.00	0.00	0.86	172.15
Subsidy Granted (EUR K)	78.62	100.55	0.00	0.00	0.00	0.49	122.38
Educ. Years	11.68	0.98	11.56	1.04	10.50	11.67	12.63
College Share (%)	15.24	16.84	15.39	18.45	0.00	12.50	34.62
Production Worker Share (%)	70.96	21.53	68.43	25.11	37.50	72.73	100.00
Observations	1600		1600			3200	

Notes: All variables measured at $\tau = -3$. Back to Section 1.4.

Table A.37: The Matched Control Group: Balance Table B (Losers vs. Matched Control).

	Treatment Group		**Control Group**		**Both**		
Variable	Mean	Std. Dev.	Mean	Std. Dev.	10p	Median	90p
Revenue (EUR M)	1.62	5.52	1.27	2.71	0.10	0.43	2.71
Employment	9.02	18.56	8.81	15.12	1.00	3.90	20.00
Wages (EUR K)	17.81	7.95	18.01	8.79	5.50	18.80	27.82
Subsidy Applied (EUR K)	47.47	76.19	0.00	0.00	0.00	0.00	65.59
Subsidy Granted (EUR K)	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Educ. Years	11.34	1.12	11.42	1.23	10.00	11.50	12.56
College Share (%)	10.50	15.47	15.41	21.76	0.00	6.90	33.33
Production Worker Share (%)	74.25	25.39	70.77	27.93	30.95	79.63	100.00
Observations	123		123			246	

Notes: All variables measured at $\tau = -3$. Back to Section 1.4.

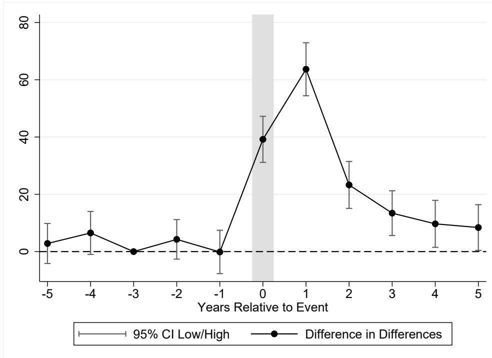

(a) Machinery Investment.

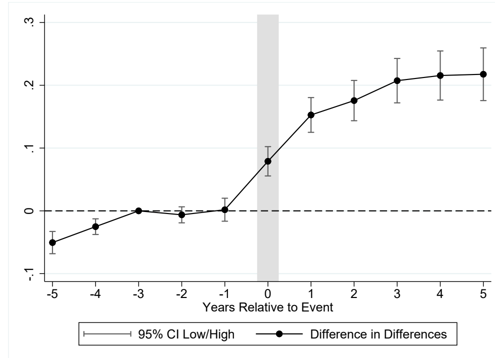

(b) Employment.

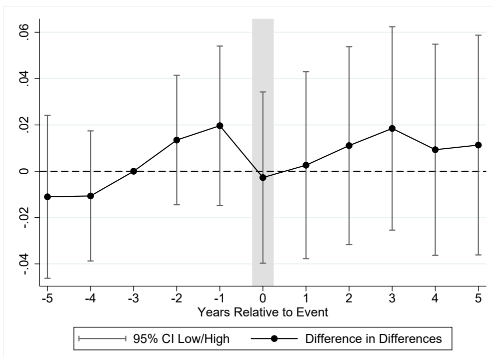

(c) Education Years.

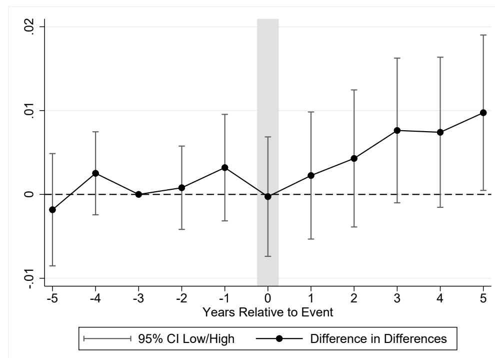

(d) College-Educated Workers' Share.

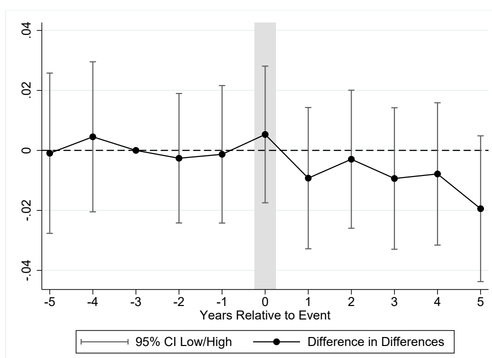

(e) Production Workers' Share.

Figure A-32: The Match Control: Event-Study Estimates.

Notes: The sample is winners matched to non-applicants (matching procedure is described in detail in the paper). The figures show event study graphs of machinery investment, employment (relative to $\tau = -3$ level), years in education, the employment share of college-educated workers, and the employment share of production workers. The treatment group is the subsidy winners (the main treatment group), and the control group is constructed via matching. We use coarsened exact matching (CEM). We match by revenue, employment, wages at $\tau = -3$ plus revenue and employment changes in percentages from $\tau = -3$ to $\tau = -1$ and industries' main sectors (letter classes). The CEM percentiles are 10, 25, 50, 75, 90, and 99. The match is 1:1 with replacement. Event time $\tau = 0$ refers to the application year. Calendar year indicators are included as controls.

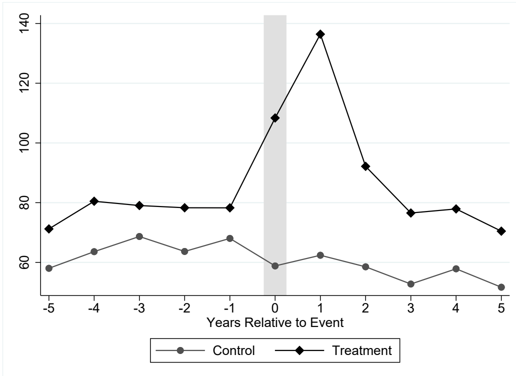

(a) Machinery Investment.

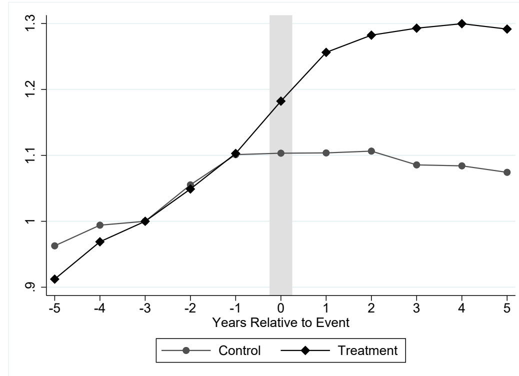

(b) Employment.

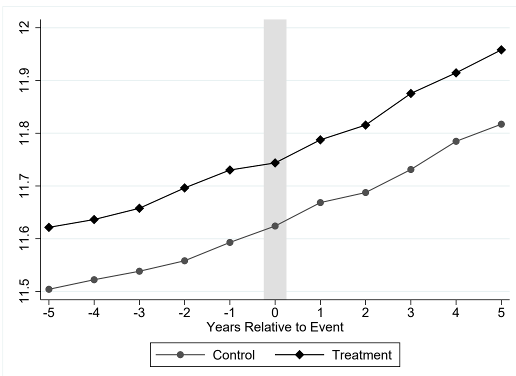

(c) Education Years.

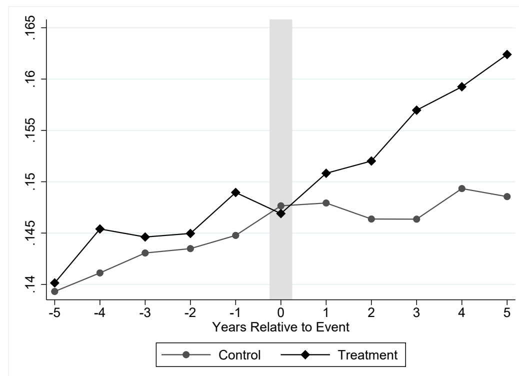

(d) College-Educated Workers' Share.

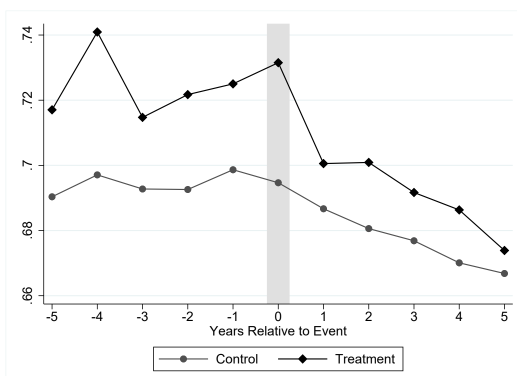

(e) Production Workers' Share.

Figure A-33: Raw Means: Winners vs. Matched Control.

Notes: The sample is winners matched to non-applicants (the matching procedure described in the paper). The figures show mean graphs of machinery investment, employment (relative to $t = -3$ level), years in education, the employment share of college-educated workers, and the employment share of production workers.

Table A.38: The Effects on Export Outcomes for the Matched Sample.

	(1) Export Status	(2) Export Share	(3) Export Regions	(4) Products	(5) Products Introduced	(6) Products Discontinued
Treatment	0.0351***	0.00707***	0.206***	0.172***	0.0686***	0.0698***
	(0.00849)	(0.00175)	(0.0245)	(0.0247)	(0.0112)	(0.0109)
Baseline	0.261	0.0466	1.353	1.622	0.527	0.578
N	3200	3200	3200	3200	3200	3200

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The estimated effects on export outcomes. The sample is the matched control sample to explore robustness of the export results on the particular sample.

A.3 The Spikes Design

To explore external validity, we consider technology adoption without the subsidy program. This design exploits the precise timing of technology investment events, which we call spikes, to analyze technologies' short-term effects at the firm level. The second design is valuable because the subsidy-based design is subject to two external validity concerns: 1) subsidy program as variation source, 2) program participants' representativeness. The spikes design complements the subsidy design by using a different variation source and a different sample. The spikes design is similar to a mass-layoff design (Jacobson et al., 1993) as it uses the precise event timing for identification and builds on the work of Hawkins et al. (2015) and Bessen et al. (2020). The design detects distinct events because technology investments tend to be temporally concentrated (e.g., Doms and Dunne 1998; Caballero and Engel 1999; Cooper et al. 1999; Nilsen and Schiantarelli 2003).

The Treatment Group We define the technology investment event, the spike, as an indicator that equals 1 when a firm's technology expenditures are significantly above average for the firm:

$$D_{jt} = \mathbf{1}\left\{\text{Technology Expenditure}_{jt} > \text{Threshold} \cdot \overline{\text{Technology Expenditure}}_{jT\notin t}\right\}$$

The average expenditure is computed over timeline T leaving out the current year t. For our main specification, we use the threshold of 4 (robust to different thresholds). We measure technology expenditure as investment in machinery and equipment from the financial statement register.

The sample design is the following. We consider years 1994–2018 and restrict the sample to manufacturing, warehouse and retail, transportation industries, and firms with full-time equivalent employees (FTE) between 10 and 750 at time $\tau = -1$ relative to the event. We focus on a balanced sample and require that the firms operate at least starting from time $\tau = -9$. With these restrictions, we can exclude new rapidly growing firms that are not relevant to our research questions and event definition and ensure comparability with the subsidies design. Very large firms tend to have several units or plants, which obscures the evaluation of the spike.

The treatment group is the firms that experience a technology investment event and satisfy the sample-design criteria. In the case of multiple spikes, we choose the largest spike and require no other spikes in window $\tau \in [-5, 8]$. Figure A-35 shows the treatment group's average technology expenditure by year. The event time is normalized around the event ($\tau = 0$). There is a clear investment spike: a significant fraction of technology investment at the firm level is associated with significant variations.

The Matched Control Group To construct the control group, we match the spiking firms to non-spiking firms. To construct a control group, we match the spiking firms to non-spiking firms. The matched control group serves as a counterfactual for what would have happened in the short term had the spiking firms not invested. We provide a theoretical basis for this comparison in Appendix A.9. We use coarsened exact matching (CEM). We match by revenue, employment, and wages at $\tau = -3$ and industries' main sectors (letter classes). The CEM percentiles are 10, 25, 50, 75, and 90. The final caliper match is the propensity score based on the same CEM variables. The match is up to 1:5 with replacement. Table A.39 shows the covariate balance for the matched samples. We match only in the pre-period cross-section to ensure that the pre-trend comparison between the treatment and control is informative.

Estimation The empirical strategy contrasts the treatment group with a spike to the matched control group that did not have a spike within the same 5-year window using a dynamic difference-in-differences design. To do so, we estimate Equations 1.1 and 1.2 from Section 1.4.2.

The First Stage Figure A-35 shows the first stage. The outcome is technology investment. Treatment group firms invested 2 million EUR more in technologies than the control firms in the event year. Before and after it, the groups invested similar amounts and were on parallel trends.

Variation We outline a theoretical framework that clarifies the source of variation in Appendix A.9, adapted from Cooper et al. (1999). The same model provides the basis also for the subsidies design, and we refer to it in Section 1.4.1. The main result of the model is that with adjustment costs, firms may experience low technology-investment activity periods followed by bursts of investment activity. The model produces a cutoff rule for the firm's optimal policy, where the firm adopts the technology if and only if the propensity $H \geq H^*$ for a cutoff H^* (Figure A-51).

This result clarifies that the treatment and the matched control group could be comparable in the short run because minor initial differences may lead to significant variations in technology investment. For example, in the model, one reason a firm invests and the other similar firm does not is that they have a different replacement cycle. Our estimates from the spikes design exploit the precise timing of technology investment events.

Robustness The estimates are robust to excluding firms that simultaneously start exporting, change their management, make significant investments in buildings and property, or open a new plant before the event, and to different controls (not reported).

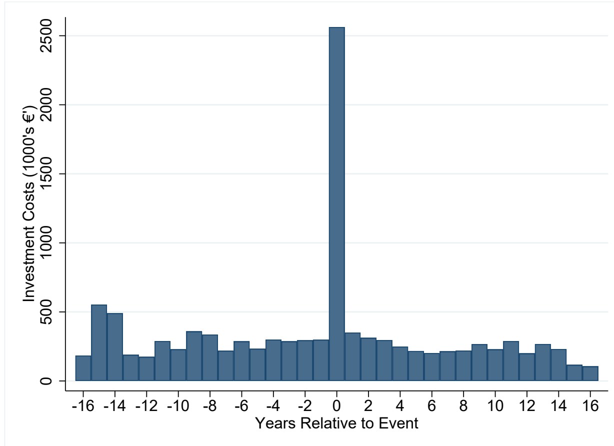

Figure A-34: The Spikes Design. Machinery Investment.

Notes: Machinery investment in EUR 1000s. Event time is normalized to zero in the year of the largest machinery investment. The sample is restricted to manufacturing, retail, transportation industries and firms with employment 10–750 for comparability with the subsidies design. Consistent with the theoretical framework in Appendix A.9, technology investment is typically a spiky activity. Back to Section A.3.

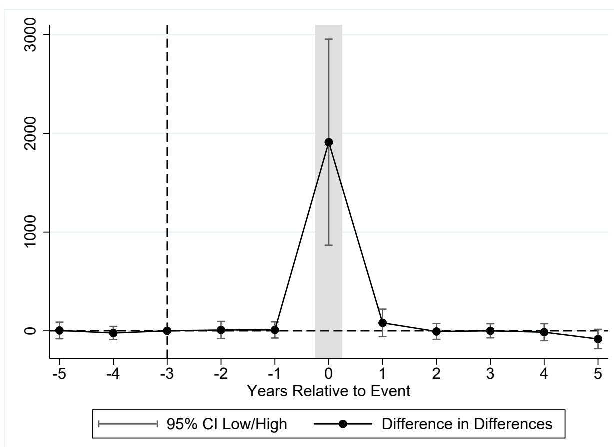

Figure A-35: The Spikes Design. First Stage: Machinery Investment.

Notes: Event-study estimates from Equation 1.1. The outcome is machinery investment in EUR 1000s. Event time is normalized to zero in the year of the largest machinery investment. Back to Section A.3.

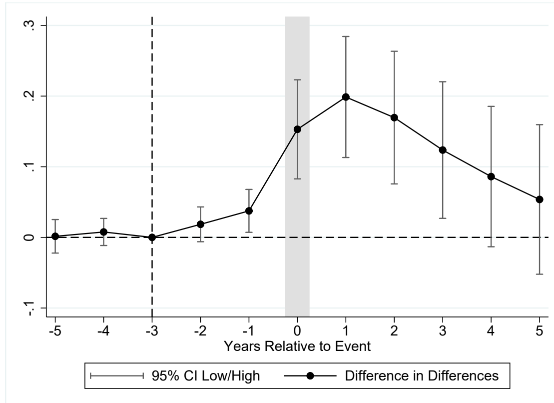

(a) Employment (%).

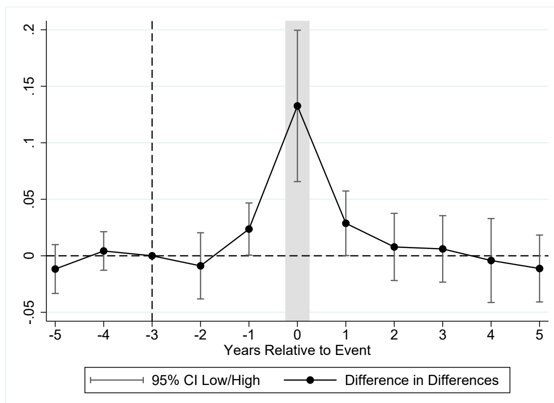

(b) Worker Entry Rate.

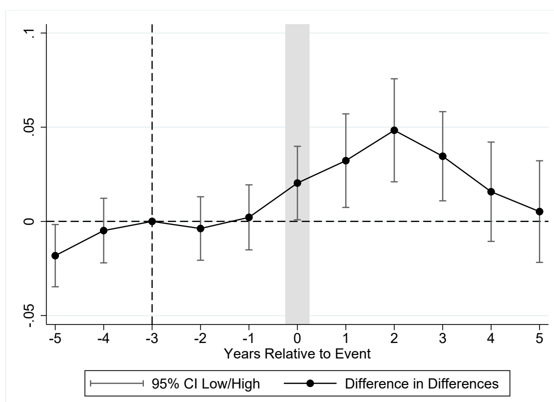

(c) Worker Exit Rate.

Figure A-36: The Spikes Design. Employment Effects.

Notes: Event-study estimates from Equation 1.1. Event time is normalized to zero in the year of the largest machinery investment. Employment is in % relative to the base year $\tau = -3$. Entry rate is defined as the number of entering workers divided by employment in the base year $\tau = -3$. Exit rate is defined as the number of exiting workers divided by employment in the base year. Back to Section A.3.

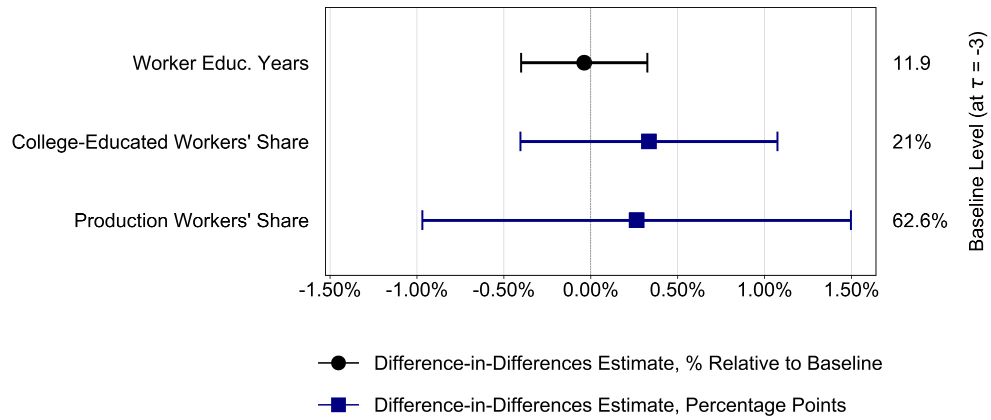

Figure A-37: The Spikes Design: Skill Effects.

Notes: Difference-in-differences estimates from Equation 1.2. The estimates compare the spikes treatment group to a matched control group. The right-hand side reports outcome means at $\tau = -3$. Back to Section A.3.

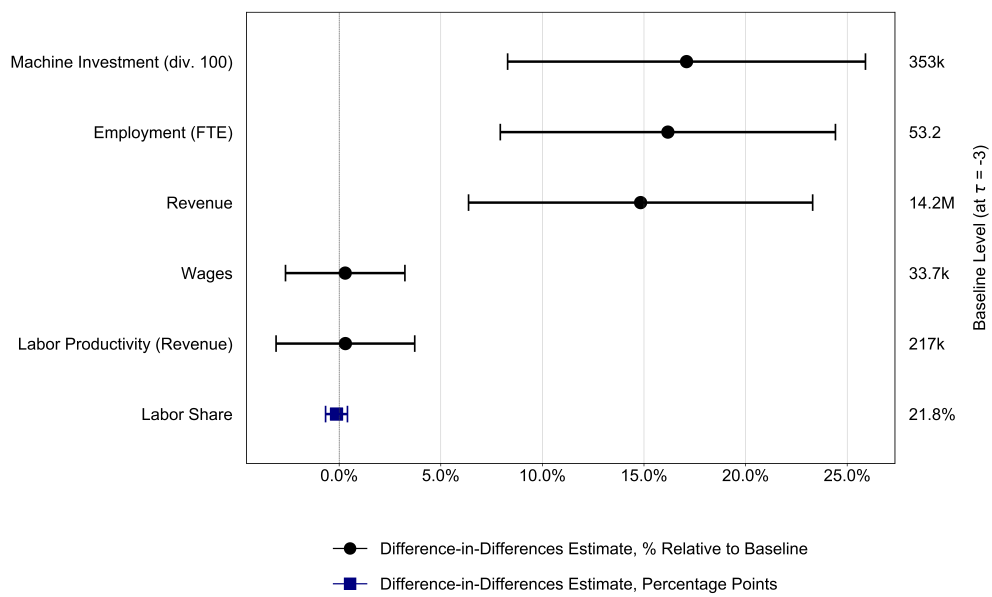

Figure A-38: The Spikes Design: Firm-Level Effects.

Notes: Difference-in-differences estimates from Equation 1.2. The estimates compare the spikes treatment group to a matched control group. The right-hand side reports outcome means at $\tau = -3$. Back to Section A.3.

Table A.39: The Spikes Design: Balance Table.

	Treatment Group		Control Group		Both		
Variable	Mean	Std. Dev.	Mean	Std. Dev.	10p	Median	90p
Machinery Inv. (EUR K)	271.21	858.93	376.70	999.26	6.96	109.55	770.54
Revenue (EUR M)	14.48	30.10	14.12	69.98	1.29	4.85	26.97
Employment	51.66	68.29	53.67	71.11	11.10	28.30	119.20
Wages (EUR K)	33.68	9.26	33.75	8.21	25.12	32.56	43.20
Subsidy Applied (EUR K)	72.40	339.62	25.23	119.99	0.00	0.00	45.37
Subsidy Granted (EUR K)	41.15	173.02	16.07	81.71	0.00	0.00	23.17
Educ. Years	11.89	0.91	11.86	0.87	10.88	11.78	12.94
College Share (%)	21.24	16.70	20.90	14.95	5.56	17.65	40.91
Production Worker Share (%)	58.90	30.95	63.68	25.79	14.29	71.43	88.89
Observations	450		1593		2043		

Notes: All variables measured at $\tau = -3$ relative to the event. We use coarsened exact matching (CEM) with replacement. Back to Section A.3.

Table A.40: The Effects by Firm Size for Spikes Design.

Panel A: Large Firms.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Sh.	(7) Educ. Years	(8) College Sh.	(9) Prod. Work. Sh.
Treatment	7235.1^{**}	0.325^{*}	0.240^{**}	0.00368	-0.00122	-0.000446	-0.0372	-0.00826	0.0284
	(2430.8)	(0.153)	(0.0893)	(0.0174)	(0.0343)	(0.00504)	(0.0412)	(0.00847)	(0.0146)
Obs.	368	368	368	368	368	368	368	368	345

Panel B: Medium-Sized Firms.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Sh.	(7) Educ. Years	(8) College Sh.	(9) Prod. Work. Sh.
Treatment	1015.6^{***}	0.0681^{*}	0.0981^{*}	0.0400	0.0141	0.000103	0.0136	0.00387	0.00949
	(149.4)	(0.0319)	(0.0411)	(0.0340)	(0.0261)	(0.00393)	(0.0274)	(0.00466)	(0.00707)
Obs.	788	788	788	788	788	788	788	788	727

Panel C: Small Firms.

	(1) Machine Inv.	(2) Employment	(3) Revenue	(4) Wages	(5) Productivity	(6) Labor Sh.	(7) Educ. Years	(8) College Sh.	(9) Prod. Work. Sh.
Treatment	391.3^{***}	0.156^{*}	0.160	-0.0256	0.00889	-0.00359	-0.00782	0.00970	-0.00517
	(46.69)	(0.0667)	(0.0895)	(0.0157)	(0.0297)	(0.00492)	(0.0415)	(0.00680)	(0.00852)
Obs.	887	887	887	887	887	887	887	887	830

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The sample is the spikes design sample. Estimated effects on selected outcomes for different firm sizes. Large firms (FTE > 75), Medium Firms (FTE >= 25 & FTE <= 75), Small Firms (FTE < 25). Machine investment is in EUR K. The effects are qualitatively similar in firms of all sizes.

A.4 The Regression Discontinuity Design

Design We use a regression discontinuity (RD) design generated by a change in the rules used to evaluate the applications as one tool to address internal validity. Buri (2017) discusses the policy change and the RD strategy. The advantage of the RD design is that the estimates are likely to reflect a causal relationship and satisfy Assumption 1. The disadvantages of the RD design in this context are statistical power, that the treatment is less precisely defined, and that it does not allow a natural way to use the text data to measure different types and uses of technology.

The EU expanded the definition of a small firm in 2005. Our RD design uses the fact that firms just below the new threshold were prioritized for subsidies but were otherwise similar to those just above it. Before the policy change, upper thresholds for small firms were 50 for employment, EUR 5M for the balance sheet, and 7M for turnover. The EU raised the thresholds for balance sheet and turnover to 10M. We use the balance sheet's total value as our running variable because it measured most precisely and had the most significant change; this gives us the statistical power to conduct the analysis.

The critical part is that the new rule was applied using retrospective data for firms. Thus firms could not immediately manipulate their size. However, as shown in Figure A-39, firms adjusted their size later. This evidence leads us to focus only on the first year of the policy change when manipulation at the threshold was unlikely. Finland implemented the change in 2007 but considered retrospective data from 2004–2006. Our estimates use 2004 data as the running variable to avoid selection bias.

The policy change potentially affected firms' self-selection into the program, the likelihood of winning the subsidy, and the levels of subsidies. While being a small firm is not a strict criterion for receiving subsidies, the ELY Centers prioritize small firms (e.g., Takalo et al. 2013). The firms know this and are potentially more likely to apply for subsidies when the expected benefits are more significant. These facts and statistical precision lead us to focus on the reduced-form effects. There were no simultaneous policy changes at the same margin.

To produce the RD estimates, we use the following specification:

$$Y_i = \alpha + \beta E_i + f(z_{i,2004}) + \varepsilon_i \tag{A.1}$$

where Y_i is outcome for firm i, $f(z_{i,2004})$ is a function of the running variable (balance sheet in 2004) and E_i is the cut-off indicator (balance sheet under 10M in 2004). We use the bandwidth of 5 million, triangular kernel, and first-order polynomial (Gelman and Imbens, 2019) in our main specification. We cluster the standard errors at the 3-digit industry.

Results Table A.42 shows the summary statistics for the RD sample firms.[1] As expected, the RD sample firms are larger than in the main design because, by definition, their revenue is around EUR

[1] We exclude agriculture and forestry, the public sector, transportation, and finance since these sectors are generally not eligible for these ELY Center subsidies.

10M. Figure A-39 documents firms starting to bunch around the new threshold after the change comes into effect. Figure A-40 formally shows by a McCrary test (McCrary 2008; implemented as in Cattaneo et al. 2018) that this is not yet the case in the pre-change year of 2004, which is the relevant year for our identification. Table A.42 tests whether firms are different on different sides of the cutoff before the treatment and finds no statistically significant differences.

Next, we describe the first stage. Figure A-41 shows a jump in the received subsidies at the new cutoff of EUR 10M. The running variable (x-axis) is the balance sheet in 2004; the outcome variable (y-axis) is the total received subsidies in EUR 10K The received subsidies are larger on the left side of the cutoff, likely because those firms became small under the new classification. Figure A-41 also shows that these subsidies stimulated new investments: The linear graphs show a clear jump at the cutoff. Table A.43 quantifies the same jumps using Equation A.1 for subsidies received and investments made in 2007. Becoming a small firm increased the subsidies by EUR 38K and investments by EUR 188K. Both estimates are significant at the 5% level.

Table A.44 presents the primary outcomes of the RD design. These results broadly confirm our main results of firm growth in employment and revenue but no skill bias. Being re-classified as a small firm increased employment by 9% and revenue by 25%. We see no changes in average wages, years of education, or the share of college-educated workers or production workers. The estimation is done by setting the average of 2003–2006 as a baseline value and comparing each observation from 2010 to 2015 separately to the baseline to increase statistical power. These differences are the outcomes in the estimation. Figure A-42 visualizes a similar estimation for each year separately. We observe an increase of 8–10 employees from 2010 onwards.

We run multiple robustness and placebo tests for our estimates. Figure A-43 explores robustness to the choice of bandwidth: Our results are not sensitive to it. Figure A-44 runs our main specification with different thresholds: We cannot replicate our results with the placebo thresholds. Figure A-45 run the estimation with placebo years' balance sheets: We observe no effect.

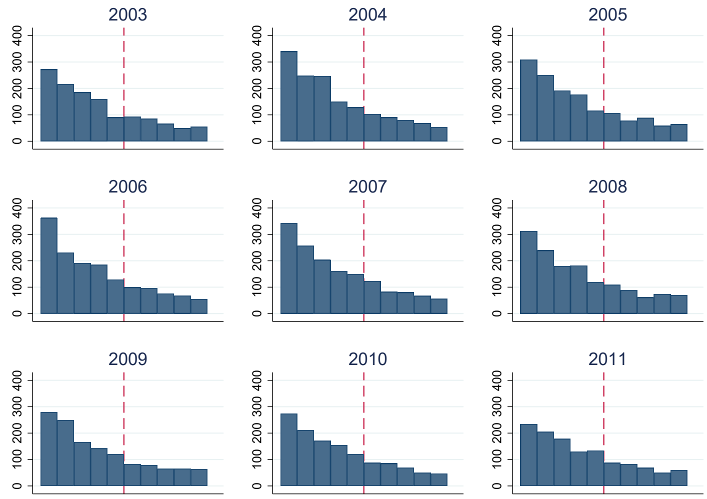

Figure A-39: RD: The Number of Firms at the Balance-Sheet Threshold.

Notes: This figure shows the number of firms around the balance-sheet threshold for small firms announced in 2003, which came into effect in 2007. Back to Section A.4.

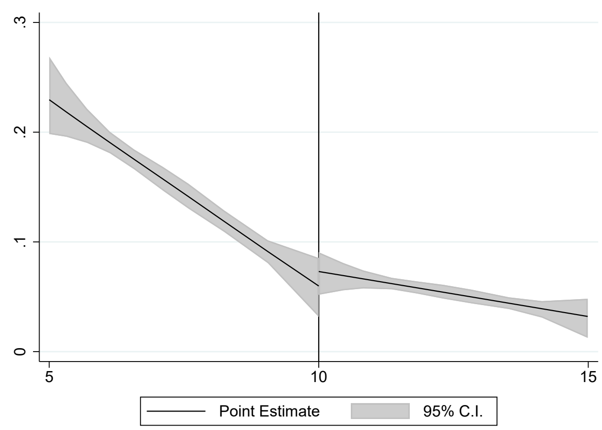

Figure A-40: RD: The Density of Firms at the Balance-Sheet Threshold.

Notes: This figure visualizes the McCrary-test for our RD year. The horizontal axis is the firms' balance sheet in 2004 in millions of euros. The vertical axis denotes the density of observations. Back to Section A.4.

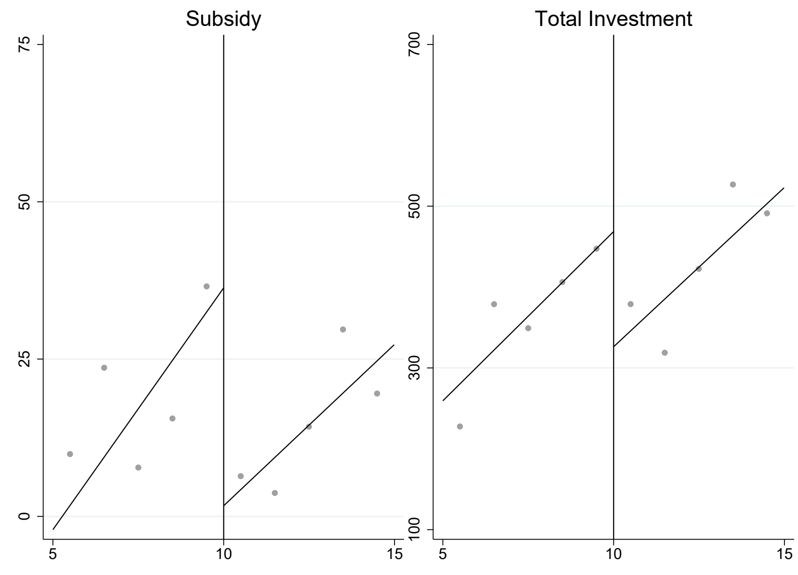

Figure A-41: RD: The First Stage.

Notes: This figure shows the discontinuity at the balance-sheet threshold for 2007 investment subsidies (left) and total investment (right). The vertical axis is in thousands of euros, and the horizontal axis is in millions of euros. Back to Section A.4.

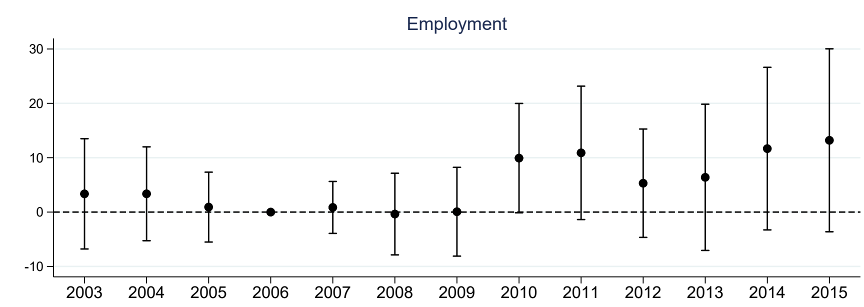

Figure A-42: RD: Employment.

Notes: The estimates are from Equation A.1. The outcome is the employment difference to base year 2006. The explanatory variable is the balance-sheet RD threshold indicator. In all regressions, we cluster the standard errors by three-digit industry, the kernel function is triangular, and the polynomial order is one. Back to Section A.4.

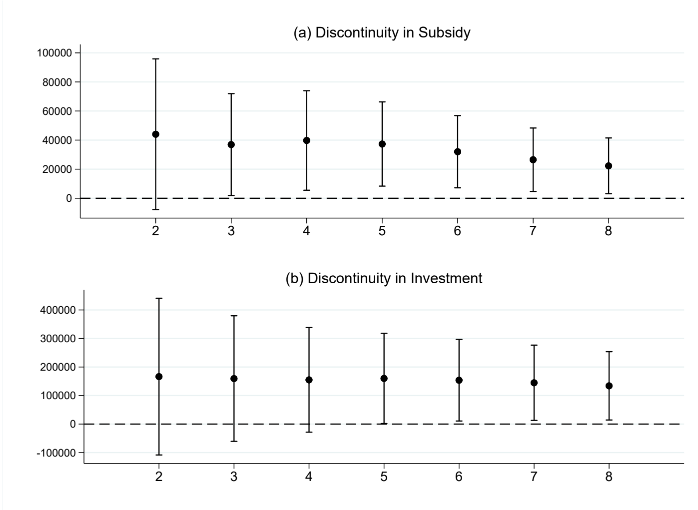

Figure A-43: RD: Different Bandwidths.

Notes: The estimates are from Equation A.1. The horizontal axis indicates the size of the estimation window. In all regressions, we cluster the standard errors by three-digit industry, the kernel function is triangular, and the polynomial order is one. Back to Section A.4.

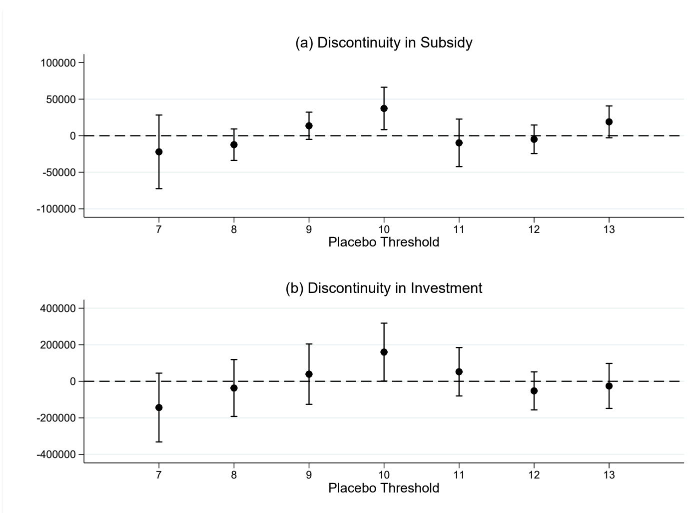

Figure A-44: RD: Placebo Thresholds.

Notes: The estimates are from Equation A.1. The outcome is investment subsidies in the upper panel and investment in the lower panel. The explanatory variable is the balance-sheet threshold indicator. The indicator equals one if the balance sheet is lower than the number indicated on the horizontal axis. The effect should be at the real threshold of 10. In all regressions, we cluster the standard errors by three-digit industry, the kernel function is triangular, and the polynomial order is one. Back to Section A.4.

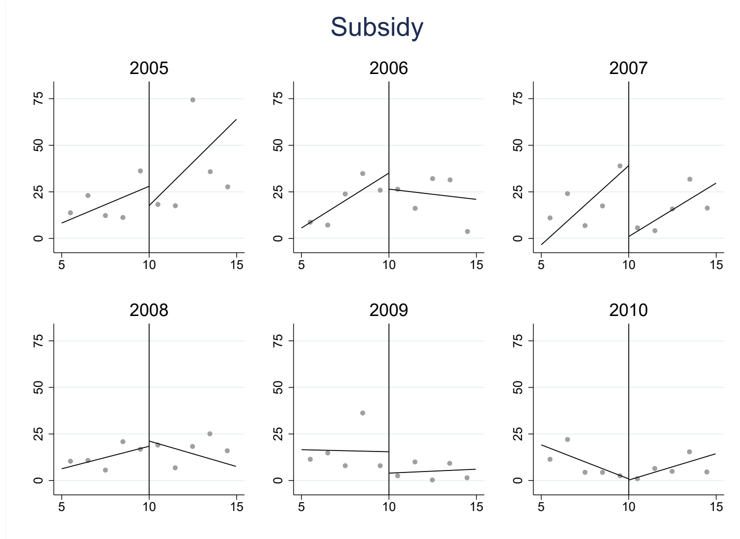

(a) Subsidies.

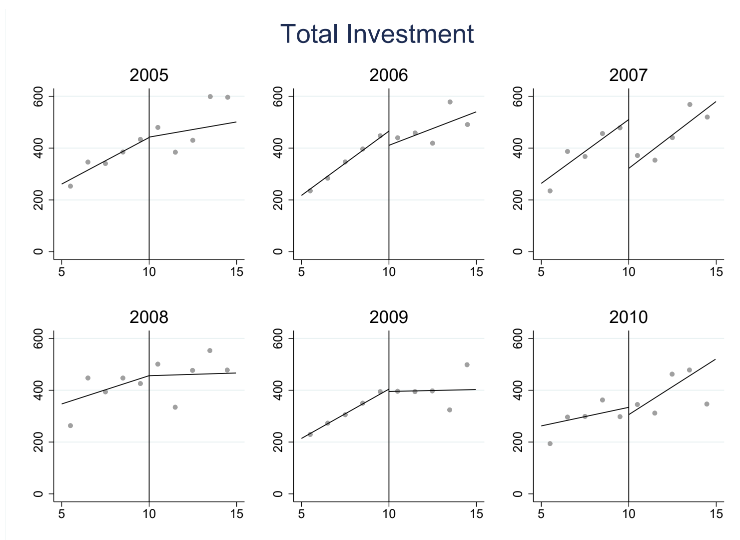

(b) Investment.

Figure A-45: RD: Placebo Years.

Notes: This figure shows the discontinuity at the balance-sheet threshold for investment subsidies (top) and total investment (bottom). The vertical axis is in thousands of euros, and the horizontal axis is in millions of euros. In all versions, we consider the 2004 balance sheet. The discontinuity should be exactly in 2007. Before 2007, there should not be a discontinuity since the new balance-sheet criterion was not yet in place. After 2007, there should not be a discontinuity since the balance sheet 2004 value was no longer relevant. Back to Section A.4.

Table A.41: RD: Summary Statistics.

	Mean	Std. Dev	N
Employment	65.75	76.93	1269
Revenue (EUR M)	16.7	16.5	1273
Wages	34,700	16,900	1269
Production Worker Share	0.40	0.32	1271
College Share	0.37	0.26	1273
Total Investment	377,600	579,000	1273
Investment Subsidies	16,200	127,600	1273
Total Subsidies	23,900	124,600	1273
Subsidized Loans	168,500	1,055,500	1273

Notes: Summary statistics for the RD sample, with balance sheet between 5 to 15 million EUR. Back to Section A.4.

Table A.42: RD: Pre-Treatment Covariate Balance.

	Investment	Subsidy	Revenue	Employment
Small 2004	5.771	16.17	-4.296	-7.745
	(88.22)	(19.03)	(2.849)	(10.37)
N	1273	1273	1273	1270

Notes: The estimates are from Equation A.1. The outcomes are pre-period averages over years 2000–2004. Standard errors in parentheses, clustered by three-digit industry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Back to Section A.4.

Table A.43: RD: The First Stage.

	(1)	(2)
	Subsidy	Investment
Small 2004	38.07*	188.5*
	(16.44)	(86.53)
N	1273	1273

Notes: The estimates are from Equation A.1. The outcomes are 2007 investment subsidies (left) and 2007 total investment (right). The values are in EUR K. Standard errors in parentheses, clustered by three-digit industry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Back to Section A.4.

Table A.44: RD: The Reduced-Form Estimates.

	(1) Employment	(2) Revenue	(3) Wages	(4) College Share	(5) Educ. Years	(6) Producion Worker Share
Small 2004	0.0899*	0.251***	0.0214	-0.00108	-0.00902	0.00613
	(0.0417)	(0.0435)	(0.0208)	(0.0106)	(0.0625)	(0.0119)
N	6005	6006	6003	6012	6012	6012

Notes: The estimates are from Equation A.1. The outcomes are defined in first differences. Standard errors in parentheses, clustered by three-digit industry. * p < 0.10, ** p < 0.05, *** p < 0.01. Back to Section A.4.

A.5 Data and Fieldwork

A.5.1 Data on Technologies

This section reports details on the technology categories primarily based on the text data.

A.5.1.1 Uses of Technologies

Process This category contains cases where the firm intends to use the technology to produce the same output type. The use of technologies to automate processes or increase automation in production is part of this category. Typical descriptions: an investment that makes operations more efficient, a productivity-enhancing investment, an investment that increases automation. These descriptions often include details, for example, which part of the production the firm intended to make more efficient. Some applications describe these advances as "solving bottlenecks," complementary to the other elements in the production.

Product This category contains cases where the firm intends to use the technology to produce a new output type. Typical descriptions: diversification of production, e.g., a new product, a new service, or a more comprehensive selection of services; improved production capabilities, e.g., the ability to work with or to manufacture larger items (very common), development of product features, such as increasing quality or the degree of processing, and transitioning to more environmentally sustainable production. This category also contains cases where the firm intends to use the technologies to expand or grow, as most of these cases also explicitly include a description of new types of customers, new output, or new capabilities.

Product and process are two opposites as to whether the improvement is within or between varieties. If the text does not specify the use of the technology on this margin, we code it as NA. Typical NA cases only specify the technology (e.g., a CNC machine) or provide limited information.

A.5.1.2 Types of Technologies

Automated vs. non-automated This category classifies cases where the technology requires no active user (automated) vs. an active user (non-automated). The classification is based on the specific technology or machinery described in the text and customs data. Automated machinery includes robots, CNC machines, automated conveyor belts, automated welding tools, etc. Non-automated machinery includes not explicitly automated machinery, for example, hand-operated tools, non-automatic welding tools, hydraulic presses, non-automatic machine tools, cutting machines, lifting equipment, pumps, furnaces, and sprayers.

Hardware vs. software This category classifies cases where the technology is physical (hardware) or not physical (software). Typical hardware includes CNC machines, welding robots, laser cutters, bending presses, surface-treatment technologies, robot arms, conveyor belts, sensors, and measurement devices. Typical software includes enterprise resource planning (ERP), computer-aided design (CAD), and production-control software.

Table A.45: Summary Statistics: Text-Category Predictions using SVM.

Class	Precision	Recall	F1-score	Test Support	Number of Cases
Not Technology (0)	0.97	0.96	0.96	1550	31022
Technology (1)	0.88	0.92	0.90	571	11887
Accuracy			0.95	2121	42909
Balanced Accuracy			0.94	2121	42909
Macro Avg.	0.93	0.94	0.93	2121	42909
Weighted Avg.	0.95	0.95	0.95	2121	42909

Notes: Test Support refers to the 10% random out-of-sample of the applications classified by hand, from which accuracy measures are computed. The number of cases refers to the total number of subsidy applications with labels (both classified by hand and predicted). Precision is the ratio of correctly predicted positive observations to the total predicted positive observations. Recall (Sensitivity) is the ratio of correctly predicted positive observations to all observations in the category. F1 Score is the harmonic mean of Precision and Recall. Accuracy is the ratio of correctly predicted observations to the total observations. Back to Section A.5.

Figure A-46: Predictive Features for Technology in the Text Data.

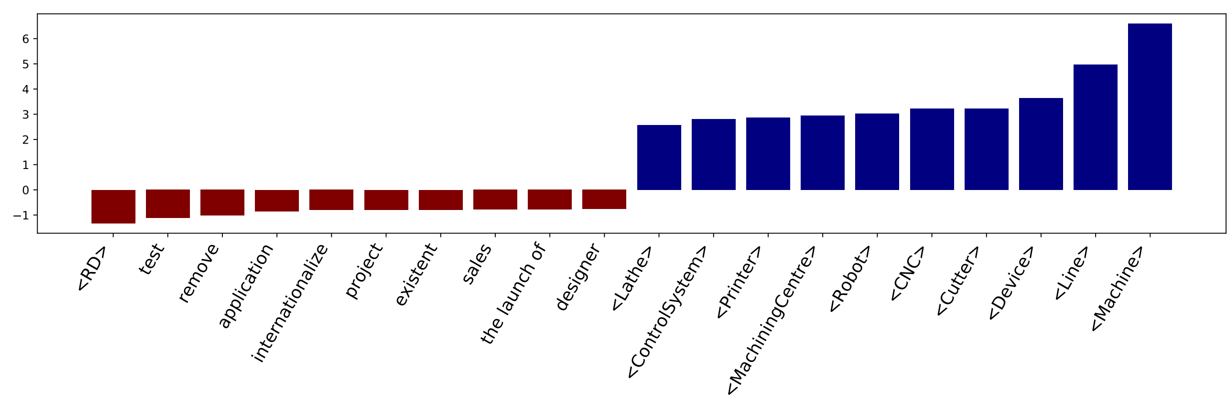

Notes: Top features for predicting technology texts. The y-axis refers to the feature weights from the SVM prediction. The features are translated into English from Finnish. Features in <> refer to compound terms combining similar spelling versions of the same term. Back to Section A.5.

A.5.2 Data on Work and Skills

We directly measure individual workers' employment, wages, education, grades, occupations, tasks, cognitive performance, and personality.

Employment and Wages We obtain employment and wage data from the registers maintained by Statistics Finland. The data contain the employment status, wages, and other income and a link to the firm. The data allow us to track all persons in Finland over time, independent of their labor-market status. The data are combined from multiple government sources (including the social security system and the tax authorities) and direct data collection by Statistics Finland. These registers also record the individuals' age and gender.

Education We measure education and school grades. Education is measured from The Register of Completed Education and Degrees. It provides exact information on the educational degrees the individual has obtained. We measure the level of education in four categories: 1) very low (no recorded degree), 2) low (high school), (3) medium (BA or equivalent), and 4) high (MA or PhD). We measure the type of education also in four categories: 1) STEM (science, technology, engineering, and mathematics), 2) HASS (humanities, arts, and social sciences), 3) business and law, and 4) other types. We map degrees to years of education based on their official length.

School grades are measured from the Secondary Education Application Register and the Finnish Matriculation Examination Board Register. We focus on the 9th-grade GPA and the standardized scores in the national high-school exit exam (12th grade).[2] We normalize both grade measures to have mean 0 and standard deviation 1 within cohorts.

Occupations and Tasks We measure occupations from the employment registers at the 3-digit level in the ISCO classification system. We harmonize the occupation classifications, resulting in 48 consistently defined occupations. For most analyses, we focus on three broad occupational categories: production workers (craft workers, operators, assemblers, and elementary occupations), non-production workers in lower-level positions (clerical, service, and sales workers), non-production workers in higher-level positions (technicians, associate professionals, professionals, and managers).

To measure the task content of the occupations, we use the European Working Conditions Survey (EWCS). The survey provides information on the tasks workers perform in their jobs. The data are collected through face-to-face interviews every five years. Using these data, we construct occupation-level measures of task intensity for routine, manual, cognitive, and social tasks (see, Autor et al. 2003).[3] For example, an occupation is classified as highly routine if the workers in that occupation describe they often perform repetitive and monotonous tasks. The advantage of the EWCS data is that it is based on workers' descriptions of their work; it is available for a specific country and time and is consistent with the European occupational classification.

[2] We use 9th-grade GPA because only approximately 50% of Finns take the high-school exit exams.
[3] We use similar classifications as Kauhanen and Riukula (2019).

Cognitive Performance and Personality We obtain data for cognitive performance and personality from the Finnish Defence Forces (FDF). The data cover 79% of Finnish men born 1962–1979, and are measured because of universal conscription. The cognitive-performance measures are visuospatial, arithmetic, and verbal reasoning. The visuospatial test is similar to Raven's Progressive Matrices (Raven and Court, 1938). The personality-trait measures are sociability, activity-energy, self-confidence, leadership motivation, achievement motivation, dutifulness, deliberation, and masculinity. The personality test is based on the Minnesota Multiphasic Personality Inventory (MMPI). We normalize all measures to have a mean 0 and standard deviation of 1 within cohorts. The FDF data are described in Izadi and Tuhkuri (2021a,b).

A.5.3 Data on Firms

We assemble a large set of data on firms, including the revenue, profits, exports, products, prices, and patents. The data track all firms over time.

Firm Performance The firm-performance measures, revenue, value-added, and profits, are obtained from the Finnish Financial Statement Register. We use two variables to measure productivity: revenue per worker and total factor productivity (TFP) estimated using the Cobb-Douglas production function. We measure profits primarily by the profit margin, defined as profits divided by the revenue. We define the labor share as the wage bill divided by the revenue. We winsorize firms' monetary values at the 5% level.

Exports Exports are measured from the Finnish Customs' Foreign Trade Statistics. We focus on the firms' export status (exporter vs. non-exporter), exports' share of the total revenue, and export destinations.

Products We measure firms' products from the Customs Register at the 6-digit CN classification. We focus on the number of products per firm and product turnover: the number of products introduced and discontinued.

Prices We compute firms' product-level prices from the Customs Register and Industrial Production Statistics. We define product-level prices as the product-level revenue divided by the number of units sold. We harmonize the product categories to be consistent over time. We focus on firm-level average prices computed as an unweighted average. We winsorize price data at the 10% level within product and year.

Patents Patent information comes from the Finnish Patent Database. We focus on the number of new patent applications per firm.

Capital We measure capital from the official records on firms' balance sheets.

Industries We measure industries at a harmonized 2-digit level classification (based on NACE Rev. 2). Our primary industry-level variable is the industry's scope for quality differentiation, which we measure using Rauch (1999), Gollop and Monahan (1991), and Sutton (1998) indices. We also measure industries' automation intensity (Acemoglu and Restrepo, 2020), tradability (Mian and Sufi, 2014) and education level (similar to Ciccone and Papaioannou, 2009).

Subsidies We measure firm subsidies from multiple registers. Two centralized systems (Yrtti 1 and 2) record the ELY Center subsidies. We gained access to these previously unstudied data, which record the application process from submission to decision. We measure other firm subsidies using the Statistics on Business Subsidies data.

A.5.4 Fieldwork

We conducted fieldwork to understand the changes we document at the level of specific firms and workers. We visited our sample manufacturing plants and interviewed CEOs, technology managers, production workers, and subsidy administrators.

Firm Visits and Interviews We chose five manufacturing firms for in-depth case studies. The primary purpose of the case studies was to observe the technologies, production, and work firsthand. We spent on average 4 hours at each manufacturing plant observing the production and conducting interviews. We also conducted five separate firm interviews (a total of 10 firms).

Our qualitative research method was open-ended interviews, building on prior qualitative research on technologies in firms (e.g., Piore 1979; Dertouzos et al. 1989; Berger 2013; Piore 2006). This method is helpful because it allows us to identify the prevalence of mechanisms we had postulated ex-ante and uncover new mechanisms that we had not anticipated. We asked the firm representatives about their production, technology adoption, motivations behind adopting technologies, the observed effects, and government subsidies.

We selected the firms to be representative of the sample and different from each other. We visited and interviewed firms with employment from 30 to 18,000 workers; subsidy winners, subsidy losers, and non-applicants; firms in rural and urban areas; privately owned and publicly traded firms; firms with high levels of own capital and firms in the corporate restructuring. All firms were in the fabricated metal product, machinery, and wood product industries.

Worker Interviews We separately interviewed five production workers using similar in-depth interviews as in our firm visits. In all interviews, we asked the respondents broadly about their work and skills, technologies they use at work, other technologies at their workplace, and the effects of technologies they had observed. Our qualitative methods draw from a long social sciences tradition to directly ask the respondents how they perceive the cause and effect. We used a semi-structured approach to interviewing that uses open-ended questions to allow a wide range of responses to emerge (see, e.g., Piore 1979; Boyd and DeLuca 2017; Bergman et al. 2019). We recruited the interview respondents in collaboration with the Finnish Industrial Union, the largest Finnish union representing industrial workers.

Subsidy Program Interviews and Text Data To understand the subsidy program, we interviewed 1) officers in all four main ELY Centers, 2) program administrators at the Ministry of Economic Affairs and Employment, 3) an external program auditor at the Ministry of Finance, and 4) a consulting firm that assists firms in subsidy applications (a total of 18 interviewees in 7 groups). We also used text records from the administrative system of the subsidy program to track the applications and qualitatively understand how the subsidy program works.[4]

[4]In addition, we studied the relevant legislature, ELY Centers' relevant strategy documents, and the official reports of the subsidy program (e.g., Ritsilä and Tokila 2005; Pietarinen 2012; Aaltonen 2013; Ramboll 2013; Auri et al. 2018; Heikkinen et al. 2019; Ilmakunnas et al. 2020, and TEM 2020).

A.6 Text Analysis

This section presents details on the text analysis performed in the paper.[5]

Preprocessing We apply similar preprocessing steps both to the description texts (used in classifying the type and use of the technology) and the evaluation texts (used in constructing the propensity scores and cosine similarity matching). These steps are:

1. All text is converted to lowercase.
2. Non-letter symbols are removed.
3. Common "stop words" are removed using the Finnish corpus in the Natural Language Toolkit (Bird et al., 2009).
4. Words are returned to their base form (also known as lemmatization) using Voikko.[6]
5. Single-character words are removed, as there are none in Finnish.
6. Words indicating firm type are removed (such as "Oy", which translates to "LLC").
7. Countries, cities, municipalities, known firm names, and technology-related words are changed to generalized versions.[7]
8. Different versions of words associated with technology are replaced with generalized versions of those words. This is mainly to generalize compound words, which are common in Finnish.[8]

Classification After the pre-processing, we turn to scikit-learn (Pedregosa et al., 2011) to perform the classification. We first transform the texts into a Bag of Words (BoW) representation, where each application text corresponds to a vector of the length of the corpus (containing all the words appearing in any texts). Then, the corresponding indices of the vector mark the number of occurrences of each word appearing in the application's text. The vectors are then transformed using term frequency-inverse document frequency (TF-IDF) weights (Salton and Buckley, 1988). The general idea of TF-IDF is to give higher weights to informative words appearing often within an application text. These weighted vectors are finally used in training a support vector machine (SVM) classifier.[9] We also performed the classification using other classifiers than SVMs, such

[5] All text-related data work is done using Python 3.7.

[6] Voikko performs a variety of NLP preprocessing tasks for Finnish text (`https://voikko.puimula.org/`).

[7] The generalized versions of the words are inside the symbols "<" and ">". For example, the word "Helsinki" (the capital of Finland) is changed to "<City>".

[8] For example, after each prior preprocessing step, the word "hitsausrobotti" (welding robot in English) is a distinct word from both "welding" and "robot". After the last preprocessing step, the word is replaced with "<Weld><Robot>" to capture its similarity to the words "<Robot>" and "<Weld>".

[9] SVMs divide the n-dimensional space of vectors (where n is equal to the length of the corpus) with a $(n-1)$-dimensional hyperplane. In the case of a binary classification problem, points on one side of the hyperplane are classified as belonging to one category, and points on the other side to the other category.

as boosting algorithms and neural network variants, but the SVMs performed the best for our purposes. To cross-validate the classifier's performance, we use K-fold cross-validation with five splits. We search the grid for optimal hyperparameters in the learning rate (or alpha), the penalty function, and several other parameters used in vectorization.[10] The score to be optimized is the F_1-score, which gives equal importance to minimizing both false negatives and false positives, as neither one is more crucial in our classification problem. The optimized parameters are (for both technology and automation classification):[11]

1. Learning rate set to .00001
2. Penalty function set to elastic net.
3. Words appearing in more than 50% of application texts are removed.
4. In addition to single words, combinations of two and three words are also used as elements in the training vectors.

This training procedure attains around 90% F-score and 95% accuracy for both technology and automation classification in our out-of-sample tests. We classify manually the applications in our sample into the remaining categories to maximize precision.

Word Vectors Word Vectors are an increasingly popular method of transforming text into numerical form to use in natural language processing tasks, as they have been shown to outperform other text presentation models in multiple different applications (Pennington et al., 2014). Word Vectors are also capable of capturing word semantics, something that simpler transformation methods are not able to do. Put simply, word vectors represent individual words as vectors, often in high dimensions. The similarity of different words can therefore be measured as the distance of their respective vectors: the closer they are in the metric space, the more similar they are.

We construct word vectors using FastText by Facebook (Bojanowski et al., 2016). FastText builds on a model by Mikolov et al. (2013) which creates vector representations of words by predicting "context words" (e.g. words appearing before or after a given word). A key feature that makes FastText attractive for our purposes is its skip-gram approach to building the word vectors: the model creates word vectors of combinations of characters also appearing inside words. That allows the model to capture better the semantic meaning of two forms of the same word. For example, the words "technology" and "technologies" both have essentially the same meaning in many contexts, but simpler models would require enough training data containing both words to construct their word vectors accurately. That is because these models treat them as entirely separate words (at least before constructing the word vectors). FastText overcomes this limitation by constructing a word vector for the common sequences of characters in both words, such as "technolog," that

[10]These include the n-gram range and the threshold for corpus specific stop words.

[11]All other parameters are set to default ones in the SGDClassifier estimator in the scikit-learn library. We tested optimizing other parameters as well using randomized search, but find virtually no improvements in accuracy.

contribute to the word vector values of all words containing the same sequence. Hence, words containing a common sequence that captures most of the semantic meaning all have similar vector representations. This aspect is especially important in morphologically rich languages such as Finnish, where various case suffixes are common.

In our application, the words appearing in the description text of each application (i.e., in our corpus) are first transformed into 100-dimensional vectors. We highlight the fact that we use the corpus of subsidy application texts to train the model, rather than using pre-trained models of the Finnish dictionary, for example. The reason for this is that words appearing in the subsidy application texts are likely to hold different semantic meanings than the same words in more general contexts. After constructing the initial word vectors, each of them is weighted by the term frequency-inverse document frequency (TF-IDF). Finally, another 100-dimensional vector is built for each application text by taking the average over each of the TF-IDF weighted word vectors in the text. Hence, we end up with each firm in our main analysis sample having one "sentence vector" giving its application texts contents in 100 dimensions. We then use these sentence vectors to build propensity score measures and match recipients to non-recipients with replacement.

Propensity Score The procedure is explained in more detail in Section 4.3 of the main text. We use the CalibratedClassifierCV estimator in the scikit-learn library to calibrate the linear SVM model, as it is not by default a probabilistic classifier. The sentence vectors are used as features and the model outputs the estimated probability of the application being successful (i.e. probability of treatment assignment).

Cosine Similarity Matching The procedure is explained in more detail in Section 4.3 of the main text. Cosine similarity gives a similarity metric between two vectors. We calculate this metric for each winner-loser pair in our main analysis sample using the sentence vectors. The match is 1:1 with replacement, so we keep only the matched loser firm with the highest similarity with a given winning firm. After manual inspection of the match quality, we also discard all matched pairs where the similarity metric between the texts is less than .85, where unity reflects identical documents.

A.7 Context Details

The Economy Finland is an industrialized, small open economy and part of the EU. The GDP per capita is similar to other northern European economies such as Germany and the UK. The industry's employment share was 21% in 2019 (OECD: 23.0%, US: 20%; ILO 2021). Finnish labor costs are close to the Euro area average and the US (Eurostat 2020; BLS 2020). Labor-market flexibility is higher than OECD average: short-term contracts are typical and authorized, but regulations constrain dismissals of regular contracts (OECD 2020). Union membership is common: 70% of workers reported being union members in 2018 (Statistics Finland 2019). Sectoral bargaining agreements set wage floors, but unions do not directly negotiate about technology adoption.

Skills Education attainment in Finland is above the OECD average: 46% of adults had obtained tertiary education in 2019 (OECD 2020). Skill measures, such as PISA for school-age students and PIAAC for adult skills, rank Finland among the world's highest (PISA 2018, PIAAC 2018). Secondary vocational education is common in Finland: it enrolls 46.5% of 17-year-olds, near the European average (OECD 2017; Silliman and Virtanen 2021). Continuous training in manufacturing firms is also common: 46.3% of manufacturing workers participated in continuing vocational training courses in 2015 (Eurostat 2021).

Trends The recent economic trends in Finland, such as manufacturing employment decline (Statistics Finland 2020), job polarization (Kerr, Maczulskij, and Maliranta 2020), and regional divergence (Böckerman and Maliranta 2001) have been similar to the US. We document the relevant trends in manufacturing firms over 1994–2018 in this Section's figures. The average education level (measured in years of education) increased from 11.5 to 12.5 years. The college-educated workers' employment share increased from 18% to 24%. The production-worker employment share declined from 68% to 62%. Similar trends apply to the wage-bill shares. The average wages increased from 22,500 EUR to 35,000 EUR per year (all monetary values are in 2017 euros). The college to non-college wage ratio increased from 1.3 to 1.37, and the production vs. non-production workers' wage ratio declined from 1.08 to .90. Productivity has increased from 120K to 180K per worker.

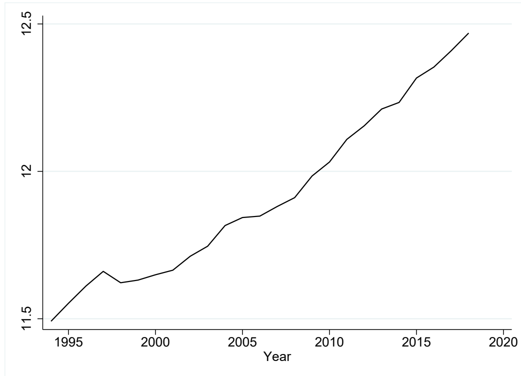

(a) Average Years of Education.

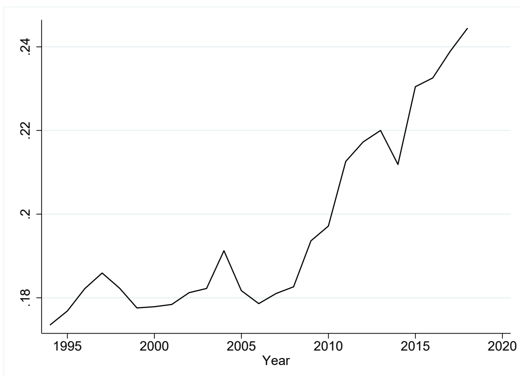

(b) College Educated Employment Share.

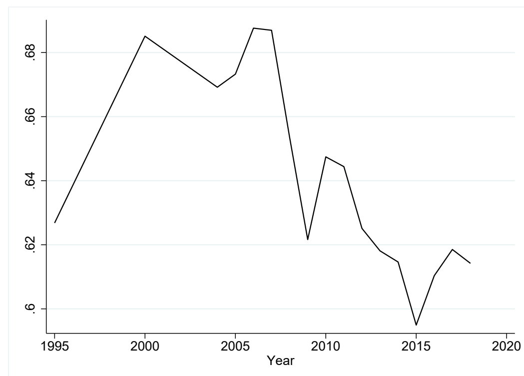

(c) Production Worker Employment Share.

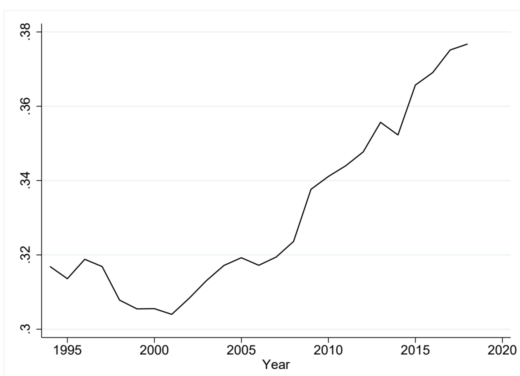

(d) College Educated Wage Bill Share.

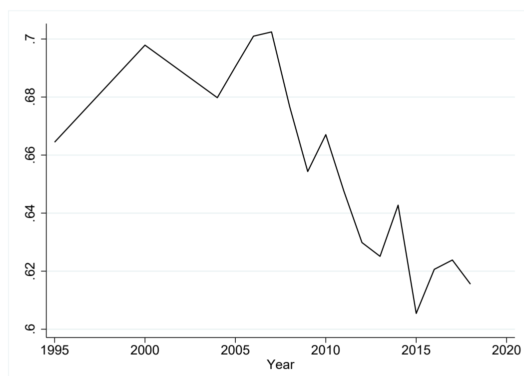

(e) Production Worker Wage Bill Share.

Figure A-47: Manufacturing Skill Composition Trends.

Notes: Trends in Finnish manufacturing over 1994–2018. We restrict to firms with at least 3 workers. We compute year-level averages from firm-level observations. The numbers are unweighted to match our research design. The employment-weighted numbers are similar.

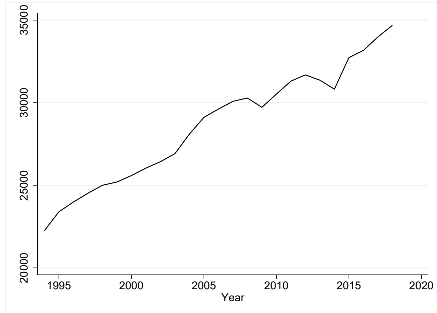

(a) Average Wages.

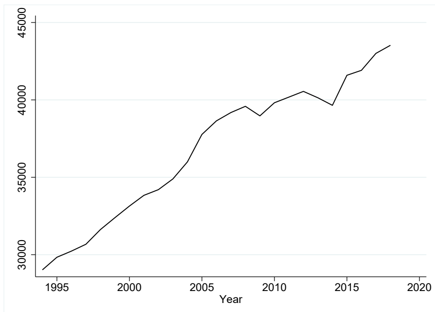

(b) College Educated Wages.

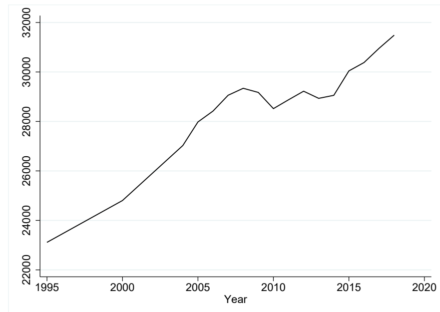

(c) Production Worker Wages.

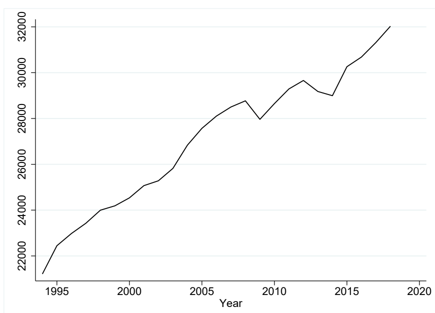

(d) Non-College Educated Wages.

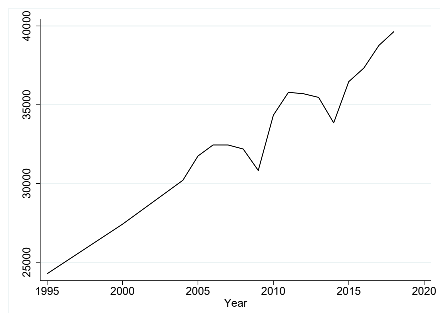

(e) Non-Production Worker Wages.

(f) College vs. Non-College Wage Ratio.

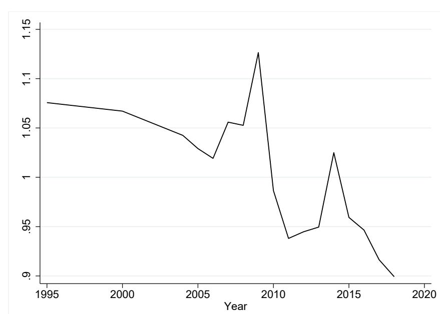

(g) Prod. vs. Non-Prod. Wage Ratio.

Figure A-48: Manufacturing Wage Trends.

Notes: Trends in Finnish manufacturing over 1994–2018. We restrict to firms with at least 3 workers. We compute year-level averages from firm-level observations. The numbers are unweighted to match our research design. The employment-weighted numbers are similar.

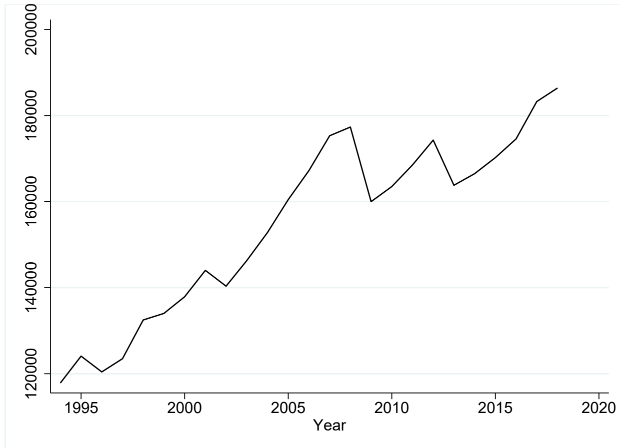

(a) Productivity (Revenue Divided by Employment).

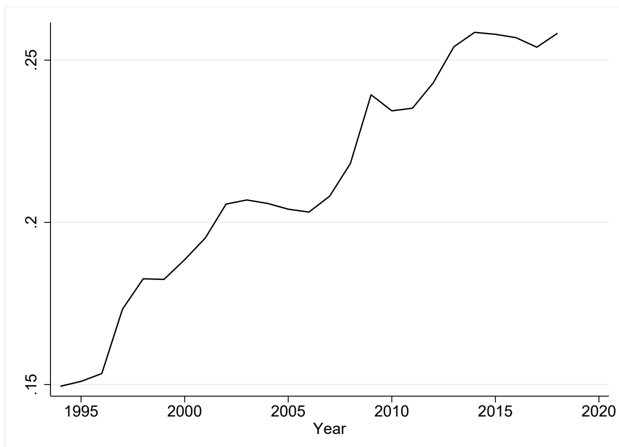

(b) Labor Share (Wage Bill Divided by Revenue).

Figure A-49: Manufacturing Productivity and Labor Share Trends.

Notes: Trends in Finnish manufacturing over 1994–2018. We restrict to firms with at least 3 workers. We compute year-level averages from firm-level observations. The numbers are unweighted to match our research design. The employment-weighted numbers are similar.

A.8 Mechanism: Predictions

In this Appendix, we collect the predictions from process and product type technological change, adapted from Melitz and Redding (2014).

A.8.1 Predictions from the Process Type

Process-type technological change has several specific and measurable implications.

Revenue Firms with lower marginal costs produce more and earn higher revenues. The CES demand structure predicts that the relative outputs and revenues of firms depend on their relative productivities:

$$\frac{q(\varphi_1)}{q(\varphi_2)} = \left(\frac{\varphi_1}{\varphi_2}\right)^{\sigma}, \quad \frac{r(\varphi_1)}{r(\varphi_2)} = \left(\frac{\varphi_1}{\varphi_2}\right)^{\sigma-1}, \quad \varphi_1, \varphi_2 > 0 \tag{A.2}$$

Productivity Lower marginal costs imply higher revenue-based productivity because of the fixed production cost:

$$\frac{r(\varphi)}{l(\varphi)} = \frac{w\sigma}{\sigma-1}\left[1 - \frac{f}{l(\varphi)}\right], \tag{A.3}$$

where input use $l(\varphi)$ is increasing in φ.

Profits Lower marginal-cost firms earn higher profits. As shown in the main text:

$$\pi(\varphi) = \frac{r(\varphi)}{\sigma} - wf = B\varphi^{\sigma-1} - wf, \quad B = \frac{(\sigma-1)^{\sigma-1}}{\sigma^{\sigma}} w^{1-\sigma} A. \tag{A.4}$$

Prices The price effect depends on whether the productivity improvement refers to lower marginal costs or a higher quality within the variety. That comes from the fact that the CES preference representation implicitly imposes a choice of units to measure the quantity of each variety. Quantity and quality are perfect substitutes within a variety, and a marginal-cost reduction is equivalent to a quality improvement, up to a new price vector. Firms with lower costs charge lower prices because the equilibrium price for each variety is a constant mark-up over marginal cost, and firms with higher quality charge higher prices because the price for each variety can equivalently be expressed in terms of quality c:

$$p(\varphi)_{cost} = \frac{\sigma}{\sigma-1}\frac{w}{\varphi}, \quad p(\varphi)_{quality} = \frac{\sigma}{\sigma-1}cw. \tag{A.5}$$

Labor Share If the composite factor of production contains only labor, the structure of the model implies that lower marginal costs reduce the labor share because the firm takes wages w as given and revenue per input increases:

$$\frac{wl(\varphi)}{r(\varphi)} = \frac{\sigma - 1}{\sigma}\left[1 - \frac{f}{l(\varphi)}\right]^{-1}. \tag{A.6}$$

If the technological change is specifically automation as in Acemoglu et al. (2020b), it substitutes capital for tasks previously performed by labor and reduces the labor share of value added.

Employment and the Labor Composition The firms use a composite factor of production L to produce. The underlying structure of the process change determines how it affects factor composition, including employment. The literature specifies different versions process-type change (Tinbergen 1975; Katz and Murphy 1992; Autor et al. 2003; Acemoglu and Restrepo 2018).[12]

In the models where technological change simultaneously reduces marginal costs and affects labor composition, technological change is typically assumed to be "skill biased," in the sense that new technologies are more complementary to high-skill workers.[13] The central prediction from these models is that if the firm adopts the technology ($T_I = 1$), the employment share of low-skill, routine, and production workers decreases:

$$s_{l^L}\left(T_I = 1\right) < s_{l^L}\left(T_I = 0\right), \text{ where } s_{l^L} = l^L / \sum_{i}^{\mathcal{L}} l^i. \tag{A.7}$$

[12]The distinction between cost and quality within the variety—while isomorphic in this framework—becomes relevant when considering the factor content of technologies. While the canonical, routine replacement, and automation models can be re-written so that instead of costs, technological change affects quality, their motivation is based on firms' cost-reduction intentions.

[13]In Autor et al. (2003) and Acemoglu and Autor (2011) the effect is mediated through tasks: technologies substitute for a set of tasks (e.g., routine or lower-complexity tasks), in which a set of workers (e.g., lower-skill workers) have a comparative advantage.

A.8.2 Predictions from the Product Type

Product-type technological change produces a set of distinct observable implications. For clarity, we consider a simplified case of two products.

Revenue Firms that introduce a new variety produce more and earn higher revenues:

$$q = \begin{cases} q(\varphi_0) \\ q(\varphi_0) + q(E[\varphi]) \end{cases} \qquad r = \begin{cases} r(\varphi_0) & \text{if } T_E = 0 \\ r(\varphi_0) + r(E[\varphi]) & \text{if } T_E = 1 \end{cases} \tag{A.8}$$

Products Firms that introduce a new variety produce a larger number of products:

$$\left|\Omega^i_{T_E=1}\right| > \left|\Omega^i_{T_E=0}\right|, \quad \omega \in \Omega \tag{A.9}$$

where $\left|\Omega^i_{T_E}\right|$ denotes the number of elements in the set of varieties produced by the firm i (measured as produced or exported products or, for example, patents).

Exports If different markets have differentiated preferences, a new variety makes it more likely that the firm starts exporting, exports a larger share of its revenue, or exports to a larger variety of destinations:

$$EXP^i_{T_E=1} > EXP^i_{T_E=0}, \tag{A.10}$$

where $EXP^i_{T_E}$ denotes the a measure of exporting activity by the firm i.

Inputs Firms that introduce a new variety use more inputs, such as labor:

$$l = \begin{cases} f + \frac{q_0}{\varphi_0} & \text{if } T_E = 0 \\ 2f + f_E + \frac{q_0}{\varphi_0} + \frac{q_1}{E[\varphi]} & \text{if } T_E = 1 \end{cases} \tag{A.11}$$

Productivity, Profits, and Prices The product-type technological change predicts, on average, zero effects on productivity, the profit margin, and prices because the expected productivity in the new variety is equal to the productivity in the existing variety. The new variety is not uniformly better than an existing variety but new and an imperfect substitute for the existing varieties. In our monopolistic-competition market structure, firms can expand either by improving productivity within a variety or by introducing a new variety, but the firms cannot expand without either action. On average, introducing a new variety appears as if the firm only scales proportionally in size. Zero effects on productivity, prices, and the profit margin combined with a positive effect on revenue are consistent with the new-varieties view.

Labor Composition, Labor Share, and Wages One critical difference between the process and product-type changes is whether technological change is likely to have distributional effects.

The product view has no unambiguous basis for expecting a sustained effect on the labor composition or the labor share. The task or skill composition might be different for the new variety, but this likely depends on the particular context.[14] The model predicts zero effects on wages in a competitive labor market because wages are determined in the sectoral equilibrium, and the firm is small relative to the market.

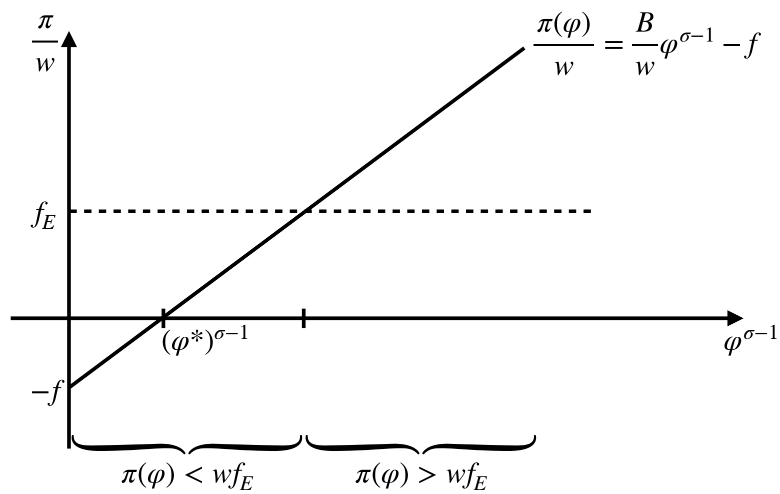

Figure A-50: The Product Cutoff.

Notes: Adapted from Melitz and Redding (2014). Back to Section 1.6.

[14]In the Nelson and Phelps (1966) view, skills are complementary to the adoption of new technologies: New technologies could induce a temporary increase in skill demand, before or after the adoption event.

A.9 Research Design: Theoretical Framework

To clarify the source of variation in our identification strategies, we consider the forces that influence a firm's technology adoption and its factor demand. We proceed in two steps. In Step 1, we focus on the firm's technology-adoption decision. In Step 2, we consider the firm's conditional factor demand, treating the technology as a quasi-fixed factor; the idea is to show that we can trace the implications of the technology adoption problem for factors' relative demand. The framework is general to allow for the analysis of multiple types of technologies and factor inputs. The adoption model is adapted from Cooper et al. (1999).

A.9.1 Step 1: Technology Adoption

In Step 1, we model the general technology-adoption problem of an individual firm. In the model, the firm makes the discrete choice between replacing existing technology with a new technology or continuing to use the old technology for another period. Consider a firm i that maximizes:

$$E_0 \sum_{t=0}^{\infty} B_t Y_t^i \tag{A.12}$$

subject to

$$Y_t^i = A_t^i \theta_t^i F\left(T_t^i; L_t^i\right) - D_t^i \Theta_t^i \tag{A.13}$$

$$T_{t+1}^i = \begin{cases} (1-\delta)T_t^i & \text{if } D_t^i = 0 \\ \tau_{t+1}^i & \text{if } D_t^i = 1 \end{cases} \tag{A.14}$$

where $\tau_{t+1}^i = \mu_t^i \tau_t^i$ and $\mu_t^i \geq 1$ is the rate of exogenous technological progress.[15] The choice variable in this problem is D_t^i where $D_t^i = 1$ if the new technology T is adopted in period t.

The first equation (A.12) is the firm's objective function. The firm maximizes the discounted present value of profits, which are defined as output minus the adjustment costs. The discount rate is $B_t \in (0,1)$.

The second equation (A.13) describes the production process and the adjustment costs. The function $F(\cdot)$ is increasing and concave in the level of technology. The output also depends on the state of productivity A_t^i. We assume that A follows a first-order Markov process $\Phi(A'|A)$. The model has two types of adoption costs. The first is a fixed adjustment cost (Θ_t^i). If the firm adopts the new technology ($D_t^i = 1$), it has to incur a cost Θ_t^i. It reflects the direct cost of the technology, its installation costs, other fixed adjustment costs, and a temporary output loss. We assume that Θ_t^i is i.i.d. The second is the opportunity cost that is proportional to the production volume. It is characterized by θ_t^i that equals $\lambda_t^i \leq 1$ during an adoption period and 1 otherwise.[16] The intuition

[15] We allow the technological progress to contain an idiosyncratic and a deterministic common component to clarify the potential mechanisms. That is, we assume $\mu_t^i = \mu_t + \varepsilon_t^i$.

[16] This implies that adjustment costs are heterogeneous across firms even if $\lambda_t^i = \lambda < 1$, i.e., equal for all firms i and periods t.

is that investment temporarily diverts resources away from production.

The third equation (A.14) describes the time path of the given technology. The technology frontier is τ_t. The firm's actual technology that is in-use is T_t^i. The in-use technology is typically less productive than the latest version because technology depreciates at an exogenous rate δ and because the latest technologies improve at rate μ_t^i. The firm can decide to adopt the latest version of the technology ($D_t^i = 1$); in that case its technology will be equal to τ_{t+1}^i in the next period. The gains to adoption reflect both technological progress (μ_t^i) and the rate of depreciation (δ).

Under this framework, the firm's technology adoption reflects several forces:

1. Replacement cycle: The underlying deterministic replacement cycle—driven by depreciation of capital δ and the common exogenous technological progress μ_t^i—will imply that the older vintage of the capital, the more likely is replacement.

2. Shocks to technologies' costs: Idiosyncratic shocks to costs Θ_t^i affect the investment in a straightforward way: lowering the costs and increasing the likelihood of the investment.

3. Shocks to technological progress: Idiosyncratic shocks to technological progress, that is shocks to μ_t^i, increase the benefits from the technology investment and increase the likelihood of the investment.[17]

4. Shocks to productivity: The response of investment to A_t^i depends on both the nature of the adjustment costs (λ_t^i and Θ_t^i) and the persistence of the shock ($\Phi(A'|A)$). The firm would prefer to replace machinery during a period where inputs are not very productive (reflecting $\lambda_t^i < 1$) and would also prefer to have a new machine available when productivity is high. To build intuition, suppose that adjustment costs are fixed. If A is i.i.d., investment is independent of A. But if a shock to A is informative of similar shocks in the future, then the investment is more likely when A is high—the firm invests now to benefit from the high productivity in the future.

We provide proofs and more detailed exposition in Section A.9.3. In the detailed version, we characterize the solution by a hazard function $H(t, A)$, the probability of adoption if the current technology stock is t and the state of productivity is A.

In words, two forces determine a technology's productivity: the technology's 'age' and a shock to total factor productivity. Given the state of productivity, the producer compares the discounted expected benefits of more productive technology relative to the current adoption costs. The gain to adoption is that a new version of the technology is more productive as it reflects some aspects of technological progress. There are two types of costs for replacement. First is the direct loss of output associated with the acquisition and installation of new capital goods. Second is that the process of installing the new machinery and retraining workers reduces productivity in the firm.

[17] Within the framework, this mechanism works analogously to the aging of technology.

The nature of the adjustment costs and the structure of the stochastic process governing the shocks jointly determine adoption timing.

The model assumes that small adjustments of technologies are either infeasible or undesirable. In particular, many technology-investment projects (e.g., the purchase of large machinery) are not possible in small quantities. In addition, the model assumes that the costs of adjusting the technologies stock may be nonconvex. Consequently, at the firm or plant level, we may see periods of low technology investment activity followed by bursts of investment activity, i.e., investment spikes. Empirical observations support this view of technology adoption: we find that a significant fraction of technology investment activity at the firm level is associated with large variations in the technology stock: i.e., technology investment is typically a lumpy activity.

A.9.2 Step 2: Conditional Factor Demand

In Step 2, we consider the firm's conditional factor demand, treating the technology as a quasi-fixed factor. This approach is closely related to the work by Berman et al. (1994) who treat machinery investments as quasi-fixed and invoke Shephard's lemma to justify their empirical specification. Cost-function estimates with quasi-fixed capital trace back to Caves et al. (1981). Our aim is to trace the implications of the technology adoption problem for factors' relative demand. The intuition is that technology is relatively more costly to adjust than labor.[18]

The firm's production function is written as:

$$Y = F\left(T; L\right) \tag{A.15}$$

where T is the technology of our focus and L is a vector of multiple other factors. An element L_i is the quantity of factor i used in the production of a quantity Y of output. We assume F is strictly increasing with each of its arguments and strictly concave. We denote the relative price of factor i by $W_i > 0$. For the purposes of this analysis, these relative prices reflect potential relative productivity effects from technology T. The conditional factor demands are characterized as solutions to the cost-minimization function:

$$\min_{(L_1 ... L_n)} \sum_{i=1}^{n} W_i L_i \quad \text{subject to} \quad F\left(T; L_1 ... L_n\right) > Y \tag{A.16}$$

The minimum value of the total cost is the cost function $C\left(W_1 ... W_n, Y\right)$. Under this framework, it satisfies the standard properties of a cost function. It is increasing, homogeneous of degree 1, and concave in $(W_1 ... W_n)$, and it satisfies the Shephard's lemma.

The Shephard's lemma gives us an analytical tool to interpret the relationship between factor demands and their prices. It states that:

$$\bar{L}_i = C_{W_i}\left(W_1 ... W_n, Y\right) \tag{A.17}$$

[18] Hamermesh 1989analyzes the costs firms face in adjusting labor demand to exogenous shocks. The study argues that adjustment costs could be viewed as fixed and documents that labor adjustment tends to be lumpy.

where $\bar{L}_i$ denotes the factor demand for the factor L_i and C_{W_i} denotes the partial derivative of the cost function C with respect to price W_i. In other words, the cost function says that the conditional factor demands can be characterized through a shock to the price vector $(W_1...W_n)$.

The expression (A.17) allows us to provide a theoretical basis for analyzing the effects of technology adoption on the demand for different types of labor. In this framework, technology's effect on labor demand is translated through its effect of the (potentially unobserved) prices of labor, which reflect the productivity of labor combined with the technology. For example, complementarity between technology and skills would mean that technology T would change the price vector $(W_1...W_n)$ in a way that the factor demands $\bar{L}_i$ would shift toward high-skill labor $L_H \in L$.

A.9.3 Details on Step 1: Technology Adoption

We consider the technology adoption (or replacement) problem of an individual firm with a given stock of technologies. This treatment is closely based on Cooper et al. (1999). The underlying technological progress in this economy makes the problem nonstationary. To analyze the problem, we normalize it to a stationary version. Define $x_t = X_t/\tau_t^i$ so that lowercase roman letters represent values which are normalized by the current value of the technology frontier. For simplicity, assume that the fixed adjustment cost is proportional to the technology frontier, i.e., $\Theta_t^i = \Theta^i \tau_t^i$ and that $F(\cdot)$ exhibits constant returns to scale. The problem is normalized as:

$$E_0 \sum_{t=0}^{\infty} \beta_t^i y_t^i \tag{A.18}$$

subject to:

$$y_t^i = A_t^i \theta_t^i t_t^i - D_t^i \Theta^i \tag{A.19}$$

$$t_t^i = \begin{cases} \rho t_t^i & \text{if } D_{t-1}^i = 0 \\ 1 & \text{if } D_{t-1}^i = 1 \end{cases} \tag{A.20}$$

In this normalized version, the discount rate (β_t^i) equals $B_t \mu_t^i$. We assume that the technological progress (μ_t^i) is not too fast so that $\beta_t^i < 1$. We define $\rho_t^i = (1-\delta)/\mu_t^i \in [0,1]$ that reflects both depreciation and obsolescence. With this normalization, technology adoption $(D_t^i = 1)$ implies that the state of the technology is 1 in the next period and a fraction ρ_t^i of its size in the previous period otherwise.

To analyze this problem, we use a dynamic programming approach. The states are the age of the technology stock (t) and the productivity shock (A). The value function $V(t, A)$ satisfies the functional equation:[19]

$$V(t, A) = \max\left[V^Y(t, A), V^N(t, A)\right] \tag{A.21}$$

where

$$\begin{aligned} V^N(t, A) &= AF(t) + \beta E_{A'|A,\varepsilon'} V(\rho t, A') \\ V^Y(t, A) &= AF(t)\lambda - \Theta + \beta E_{A'|A} V(1, A') \end{aligned} \tag{A.22}$$

[19]For expositional clarity, we drop the subscript t and the superscript i.

The superscript Y refers technology adoption ($D_t^i = 1$) and N to no technology adoption ($D_t^i = 0$). The expectation over A' is taken using the conditional distributions $\Phi(A'|A)$. We assume shock follows a first-order Markov process. The productivity shock has two effects: a direct effect on current productivity and an indirect effect through information about future productivity shocks through $\Phi(A'|A)$. We assume shocks to Θ^i are i.i.d.

The solution to the functional equation leads to adoption if and only if $V^Y > V^N$ given the state vector, $h = (t, A)$. We characterize the solution by a hazard function $H(t, A) \in [0, 1]$, the probability of adoption if the current technology stock is t and the state of productivity is A. The cutoff is visualized in Figure A-51.[20]

Proposition 1. *There exists a solution to the functional equation.*

Proof. The solution's existence follows from Theorem 9.6 in Stokey et al. (1989) if $\beta < 1$. ■

Proposition 2. *$H(t, A)$ is decreasing in t.*

Proof. For a given value of productivity A let $t^*(A)$ satisfy $V^N(t, A) = V^Y(t, A)$ where

$$V^N(t, A) \equiv At + \beta V(\rho t, A') \tag{A.23}$$

$$V^Y(t, A) \equiv At\lambda - \Theta + \beta EV(1, A') \tag{A.24}$$

Define $\Delta(t, A) = V^Y(t, A) - V^N(t, A)$. Using this object, it is sufficient to show that $\Delta(t, A)$ is decreasing in t. From (A.23) and (A.24):

$$\Delta(t, A) = At(\lambda - 1) - \Theta + \beta E_{A'}\left[V(1, A') - V(\rho t, A')\right] \tag{A.25}$$

where $V(t, A) \equiv \max\left\{V^Y(t, A), V^N(t, A)\right\}$. The first term is decreasing in t. The last part of this expression is also decreasing as t increases since $V(t, A)$ is an increasing function of t. Thus $\Delta(t, A)$ is decreasing in t. This proves that given the state of productivity A, the hazard $H(t, A)$ is decreasing in t. ■

Proposition 3. *$H(t, A)$ is decreasing in Θ.*

Proof. Using the definition of $\Delta(t, A; \Theta)$, we have

$$\Delta(t, A; \Theta) = At(\lambda - 1) - \Theta + \beta E_{A'}\left[V(1, A'; \Theta) - V(\rho t, A'; \Theta)\right] \tag{A.26}$$

The term $\Delta(t, A; \Theta)$ is decreasing in Θ and thus the result is immediate. ■

Proposition 4. *$H(t, A)$ is independent of A if $\Theta > 0$, $\lambda = 1$, and A is i.i.d.*

[20] While given the state vector, the probability of an investment spike is deterministically either zero or one, this hazard is a useful object because the idiosyncratic shocks are generally not measured in the data.

Proof. Using the definition of $\Delta(t, A)$, for the case of $\Theta > 0$ and $\lambda = 1$, we have

$$\Delta(t, A) = -\Theta + \beta E_{A'} \left[V(1, A') - V(\rho t, A') \right] \tag{A.27}$$

Since A is i.i.d., the right side is independent of the current realization of the shock. Thus the gains to replacement are independent of A. ■

Proposition 5. *$H(t, A)$ is increasing in A if $\Theta > 0$, $\lambda = 1$, and $\Phi(A'|A)$ is decreasing in A.*

Proof. Using the definition of $\Delta(t, A)$, for the case of $\Theta > 0$ and $\lambda = 1$, we have

$$\Delta(t, A) = -\Theta + \beta E_{A'|A} \left[V(1, A') - V(\rho t, A') \right] \tag{A.28}$$

The expectation over A' is conditional on A so that the current state of productivity does influence the replacement choice even though $\lambda = 1$. Since high values of A put, by assumption, more weight on high values of A', it is sufficient to show that $V(1, A) - V(t, A)$ is increasing in A for any t. This is, in turn, equivalent to the condition that

$$\int_t^1 V_{tA}(z, A) dz > 0 \tag{A.29}$$

for all t. This condition is satisfied if $V_{tA}(t, A) > 0$ for all (t, A). From (A.23) and (A.24) this positive cross-partial condition holds when $\Theta > 0$ and $\lambda = 1$. To see this, note that by assumption, replacement will eventually occur so that (A.23) is a sequence of current period returns with positive cross partials between t and A. From (A.24), $V^Y(t, A)$ has a positive cross partial since the second term is independent of t. ■

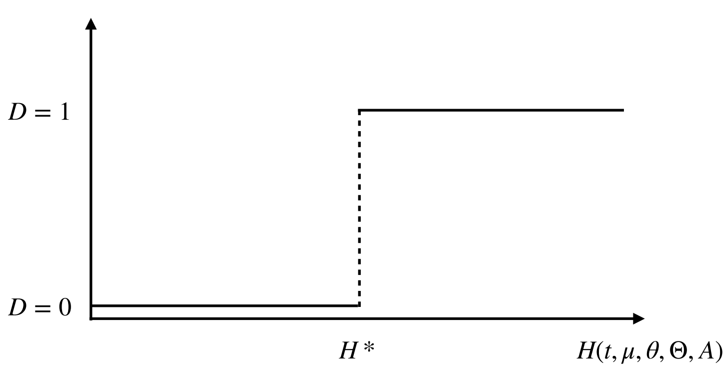

Figure A-51: The Cutoff.

Notes: Threshold model. The technology adoption model rationalizes firms' spiky investment behavior. In the model, the firm makes a technology investment $D = 1$ if adoption likelihood H crosses a threshold.
Back to Sections 1.4, A.3, and A.9.

A.10 Related Research

The Effect of Technology on Employment and Skill Demand This paper contributes to the active literature on the effects of technologies on employment and skill demand, surveyed by Acemoglu (2002b), Card and DiNardo (2002), and Acemoglu and Autor (2011), and specifically to the evidence on the effects of advanced technologies in manufacturing firms.

The closest papers to our research report similar findings. Doms et al. (1997) report little correlation between technology adoption and skill upgrading in US manufacturing, focusing on similar technologies (e.g., CNC machines and robots) and industries (e.g., fabricated metal products) as we do. Bartel et al. (2007) show that valve plants that adopted new IT-enhanced equipment shifted their business strategies toward producing more customized products, consistent with our interpretation and evidence. They report changes in machine operators' skill requirements, not in the traditional sense of replacing production workers or increasing the demand for formal education, but, for example, increased focus on setting up, monitoring, and correcting the new machinery, consistent with what we find in our fieldwork. Weaver and Osterman (2017) emphasize that most manufacturing work does not require high levels of formal education.

Additionally, Criscuolo et al. (2019) analyze the effects of an investment-support program in UK manufacturing using an instrumental variables (IV) strategy, and find evidence for a positive treatment effect on employment. Similarly, Curtis et al. (2021) find positive employment effects and no skill bias from a capital-investment tax policy in the US. Pavcnik (2003) documents that plants' adoption of foreign technology is not associated with skill upgrading, and Nilsen et al. (2009) find no evidence that investment spikes are associated with changes in the composition of the workforce. In recent work, Genz et al. (2021) report that the adoption of CNC machines and industrial robots led to increases in employment, including production workers, and did not coincide with a higher demand for more educated workers, and Koren et al. (2020) report positive wage effects on machine operators exposed to imported machinery. Extensive qualitative evidence corroborates these observations (e.g., Sohal 1996; Small 1999; Berger 2013, 2020).

Contemporary evidence on effects of robots and automation in firms supports our findings (Acemoglu et al., 2020b; Aghion et al., 2020; Bonfiglioli et al., 2020; Dixon et al., 2021; Eggleston et al., 2021; Koch et al., 2021; Stapleton and Webb, 2020). Most of it finds positive effects on employment, no negative effects on low-skill workers, and no major changes in skill composition.[21] Dixon et al. (2021) document that robot adoption is motivated by improving product and service quality, not reducing labor costs. Koch et al. (2021) report that the employment increases applied to all types of workers and provide evidence supporting the idea that exports facilitate the expansion effects of technologies. Aghion et al. (2020) report no different effects across skill groups. In contrast, Acemoglu et al. (2020b) estimate 0–1.6% declines in the production employment share while focusing on unskilled industrial jobs. The most significant difference between these studies is the result on the labor-cost share: e.g., Acemoglu et al. (2020b) and Koch et al. (2021) find labor

[21] Humlum (2019) provides a contrasting view that robot adoption affects firm-level skill composition.

share declines (3–5% and 5–7%), but Aghion et al. (2020) find no change. One way to reconcile these estimates is that the former two focus exclusively on robots, while the latter uses a broader measure of technologies. Robots specifically appear to reduce the labor share, while other advanced technologies appear to have neutral effects.[22]

Our results are different from some firm-level studies that focus on different types of technologies. Generally, the evidence suggests that investments in digital technologies may have been skill biased—in contrast to typical physical technology investments in manufacturing. For example, Akerman et al. (2015) study the regional rollout of broadband internet in Norway using a difference-in-differences design. More effective internet is a critical technological advance, but different from new manufacturing technologies, and we would expect potentially different effects. The estimates indicate that college-educated workers' wages and employment increased modestly in places that received faster internet. There were, on average, no negative effects on non-college and manual workers, but a small negative effect on high-school dropout and routine (cognitive) workers' wages. In another example, Gaggl and Wright (2017) estimate the effects of a temporary tax allowance on ICT investments, primarily software, in the UK using an RD design. They find that ICT subsidies induced increases in employment and wages. Workers performing non-routine cognitive tasks experienced the increases, routine cognitive workers experienced modest declines, and manual workers experienced no change. Furthermore, Bresnahan et al. (2002) report complementarities between skill and IT equipment, such as computers. Caroli and Van Reenen (2001) document that organizational change, Boler (2015) that R&D, and Leiponen (2005) and Lindner et al. (2021) that innovation is complementary to skills. The contrast to these papers highlights that distinct technological advances may induce distinct effects. Specifying the technologies in focus, as these papers do, is valuable for building cumulative evidence.

Our results are also different from studies that specifically focus on the replacement effects of technologies. These papers' results highlight that some technological changes may also replace workers. Bessen et al. (2020) study the effects of automation events on incumbent workers, measuring automation from firms' expenditures on third-party automation services. Our event-study design builds on their approach. The main difference is that their approach is designed to capture the replacement effects; they isolate what happens to the incumbent workers when firms automate. They find that a large increase in automation expenditure makes workers more likely to separate from the firm. The effects are meaningful but modest in size: the average earnings loss is 2%. They detect no differences by wage groups, often used as a proxy for skill. In another study, Feigenbaum and Gross (2021) analyze the replacement of telephone operators for mechanical switching by AT&T in 1920–1940. This eliminated most of these jobs, did not reduce future cohorts' overall employment, but caused adverse effects on incumbent operators.

Our results are different from several macro-level studies. We organize the macro evidence into indirect and direct approaches. The indirect approaches include Katz and Murphy (1992); Beaudry et al. (2010); Lewis (2011); Acemoglu and Restrepo (2020); Dauth et al. (2021). These papers report

[22] Similarly to Koch et al. (2021), we find zero effects on the labor share from advanced technologies.

skill bias from technological advances, partly for different reasons. The main argument in Katz and Murphy (1992) is that to reconcile the increased college wage premium with the increased supply of college-educated workers, substantial growth in the demand for more-educated workers is necessary. This demand growth is sometimes interpreted as skill-biased technological change. Similarly, Beaudry et al. (2010) and Lewis (2011) evaluate technology-skill complementarity using variation in skill supply. They find that the local skill supply predicts increases in technology adoption. This observation is consistent with our results, despite the seeming contradiction: Technology adoption may be easier in places with more high-skill workers, even if technologies do not directly affect skill composition within firms.[23] Acemoglu and Restrepo (2020) and Dauth et al. (2021) also analyze technology-skill complementarity indirectly at the local level in the US and Germany. They focus on the places' exposure to robots based on their pre-existing industry structure. This exposure approach has many clear advantages, including the possibility to analyze equilibrium effects, but the focus on variation stemming from pre-existing industries may leave out technologies' other effects than replacement, such as using technologies to launch new products.

The direct approaches include Berman et al. (1994); Autor et al. (1998); Krusell et al. (2000); Autor et al. (2003); Spitz-Oener (2006); Michaels et al. (2014), and Graetz and Michaels (2018). These papers also report skill bias from technological advances. Part of the direct macro evidence considers different technologies. Berman et al. (1994); Autor et al. (1998); Spitz-Oener (2006); Autor et al. (2003), and Michaels et al. (2014) focus on the effects of ICT, especially computers. Another part, e.g., Krusell et al. (2000) and Graetz and Michaels (2018), considers similar technologies to our study and still finds skill bias. While we do not have a complete explanation for the difference, micro and macro estimates may be different and still consistent with each other for several reasons, for example, due to externalities (see, e.g., Oberfield and Raval 2021) or if technologies induce broad economy-wide changes.[24] Exploring these channels is a promising avenue for future research.

To summarize the evidence from the prior literature, we make six observations:

1. According to contemporary evidence, technology investments in manufacturing have not appeared to cause adverse effects to workers generally.

2. Advanced technologies in manufacturing, such as CNC machines, appear to have caused increases in employment and no changes in the skill composition at the firm level.

[23] This interpretation is consistent with the technology view emphasized by Nelson and Phelps (1966); Welch (1970); Schultz (1975), where education fosters the process of technology adoption and with models of directed technological change (Acemoglu, 1998, 2002a). The interpretation is also consistent with Doms et al. (1997), who find that plants that adopted more technologies employed more educated workers *before* adoption.

[24] These reasons include: 1) externalities, e.g., in the product market, the intermediate input market, the factor market, or due to technological externalities, 2) compositional effects, e.g., through expansion and contraction of firms and industries, 3) technologies creating new areas in the economy, e.g., video games, the Apollo program, or Google, and 4) technologies directly inducing macro-level changes, e.g., self-booking platforms displacing travel agents, the internet changing job search, or technologies inducing broad organizational and cultural changes. The papers addressing externalities and compositional effects include Acemoglu et al. (2020b); Aghion et al. (2020); Humlum (2019); Koch et al. (2021); Restrepo and Hubmer (2021), and Oberfield and Raval (2021).

3. Robots, specifically, also appear to have caused increases in employment and no significant skill bias at the firm level but may have reduced the labor-cost share.

4. Digital technologies—ICT, computers, software, and the internet—appear to have been skill biased for cognitive work at the micro and macro levels.

5. Some technological advances, such as automation consulting services, appear to have caused some worker displacement.

6. Local skill supply appears to foster technology adoption.

Our results corroborate 1–3 and are consistent with 4–6. These conclusions are tentative due to the still limited evidence.

The Effects of Industrial Policy Our analysis contributes to the literature on industrial policy. By industrial policy, we refer to policies that stimulate specific economic activities and promote economic development. These policies are common. For example, EU countries spent EUR 134.6 billion on government subsidies to the private sector (designated as state aid) in 2019, about .81 % of the EU's GDP (The EU State Aid Scoreboard, 2020). The objectives and effects of industrial policy are debated (Lane, 2020).

This paper focuses on a particular type of firm subsidy: a lump-sum transfer to increase technology adoption in manufacturing. Manufacturing subsidies are widespread (see, e.g., Gruber and Johnson 2019) but understudied. Berger (2013) argues that these types of programs have contributed to the productivity and growth opportunities in German SME manufacturing, and lack of them may contribute to the relatively low productivity growth of US manufacturing. Our evidence from Finland shows that it is possible to increase technology adoption by targeted subsidies and, by doing so, induce increases in the subsidized firms' employment, revenue, and exports.

Empirical challenges in the industrial policy literature are similar to those in the literature on technology and work. There are different types of industrial policies in different contexts, and evaluating them is challenging. This paper provides new quasi-experimental estimates of firm subsidies' effects in a specific context. In addition to the research we mentioned earlier, Takalo et al. (2013) and Einio (2014) analyze Finnish R&D subsidies.

Production and Innovation Our analysis relates to the research on firms' product and export choices, intermediate inputs, and innovation.

Recent research documents that becoming an exporter stimulates technology adoption and product-quality upgrading in firms (Verhoogen, 2008; Lileeva and Trefler, 2010; Bustos, 2011; Kugler and Verhoogen, 2012). Our research finds that technology adoption also induces firms to become exporters and introduce new product varieties—the complementarity between technology and exporting appears to operate in both directions.

Access to new machinery is an example of access to new intermediate inputs. Existing research finds that access to new imported inputs fosters introducing new product varieties and productivity (Goldberg et al., 2010; Koren et al., 2020). Our research corroborates the result on product varieties. In related work, Bernard et al. (2010, 2011) analyze the role of product switching as a source of reallocation within firms, and Hausmann et al. (2007) consider product-specialization patterns' implications for growth.

Our theoretical framework builds on the literature on heterogeneous firms and trade, reviewed by Melitz and Redding (2014). We use modeling techniques from Bustos (2011) to capture the technology-adoption decisions by heterogeneous firms. We find that the monopolistic competition view of the industrial manufacturing market is consistent with our quantitative and qualitative evidence. More generally, our research provides empirical evidence to enrich the models of firm-level technological change and innovation (e.g., Hopenhayn 1992; Ericson and Pakes 1995; Klette and Kortum 2004; Acemoglu et al. 2018; Kerr et al. 2020).[25]

[25] Back to Section 1.1.

Appendix B

Appendix to Chapter 2

B.1 Supplementary Figures and Tables

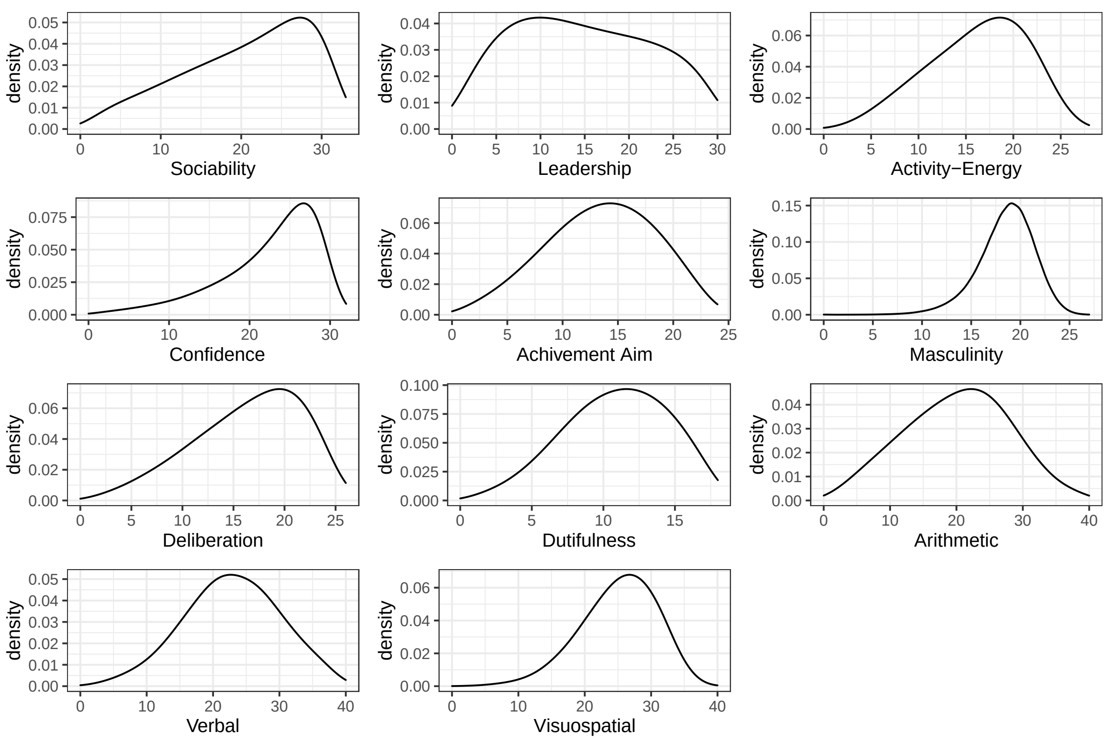

Figure B-1: Density Plots of the Raw Test Scores.

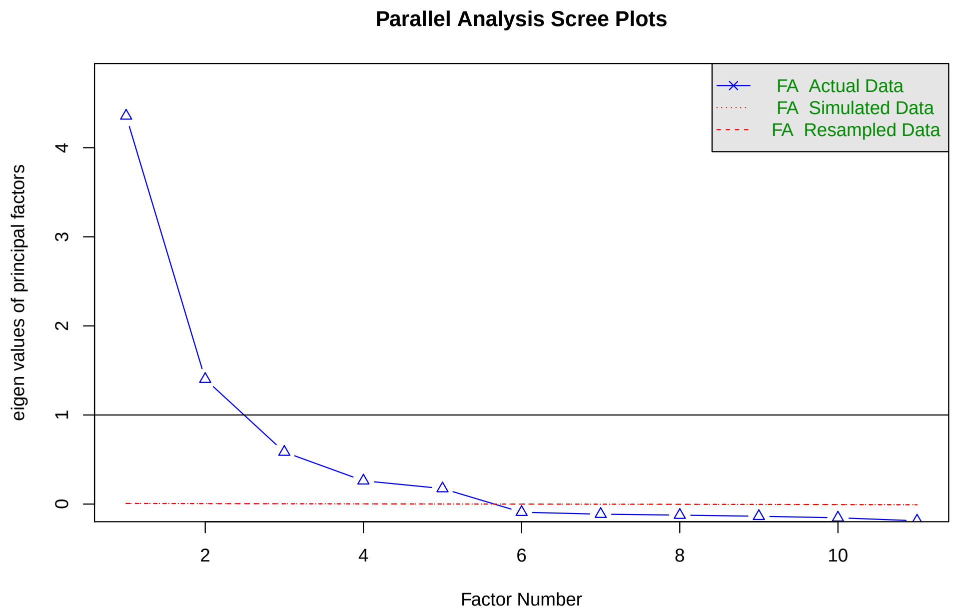

Figure B-2: Scree Plot of the Eigenvalues from Exploratory Factor Analysis of the Personality and Cognitive Test Data.

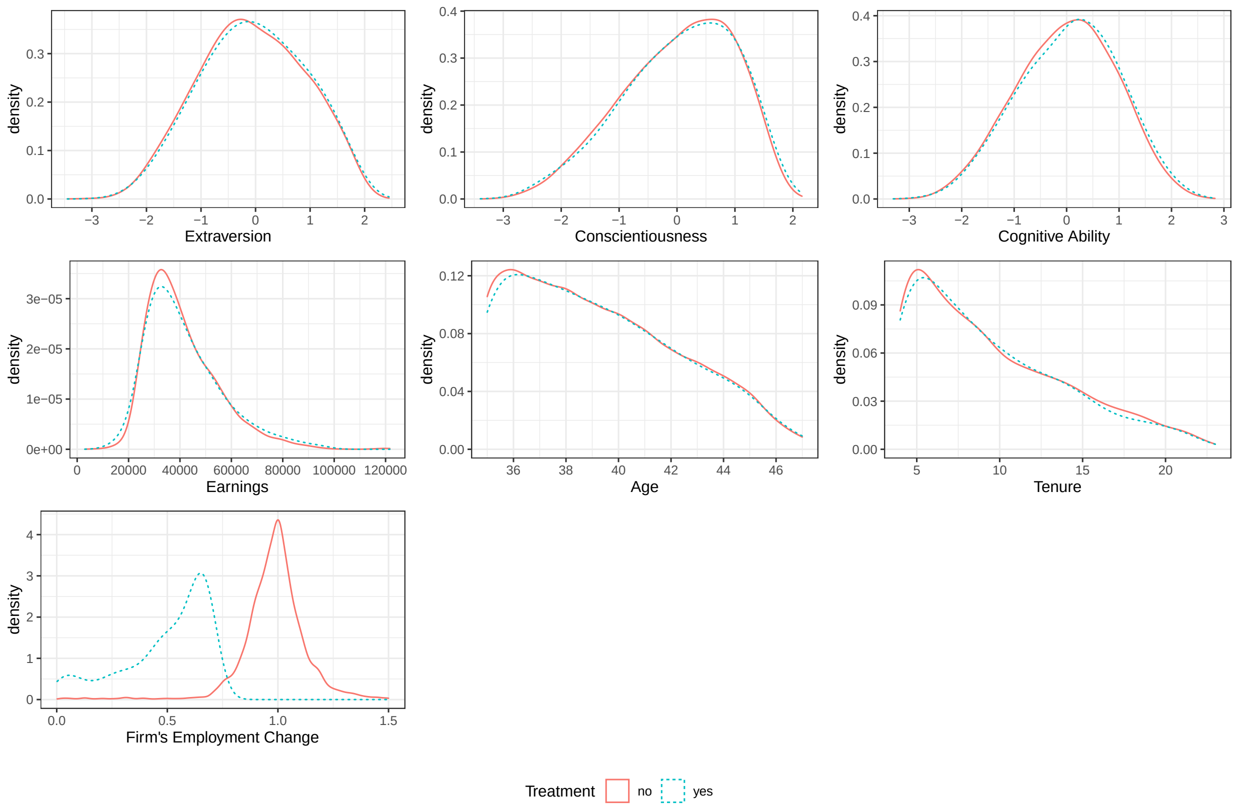

Figure B-3: Descriptive Distributions for the Mass-Layoff Sample.

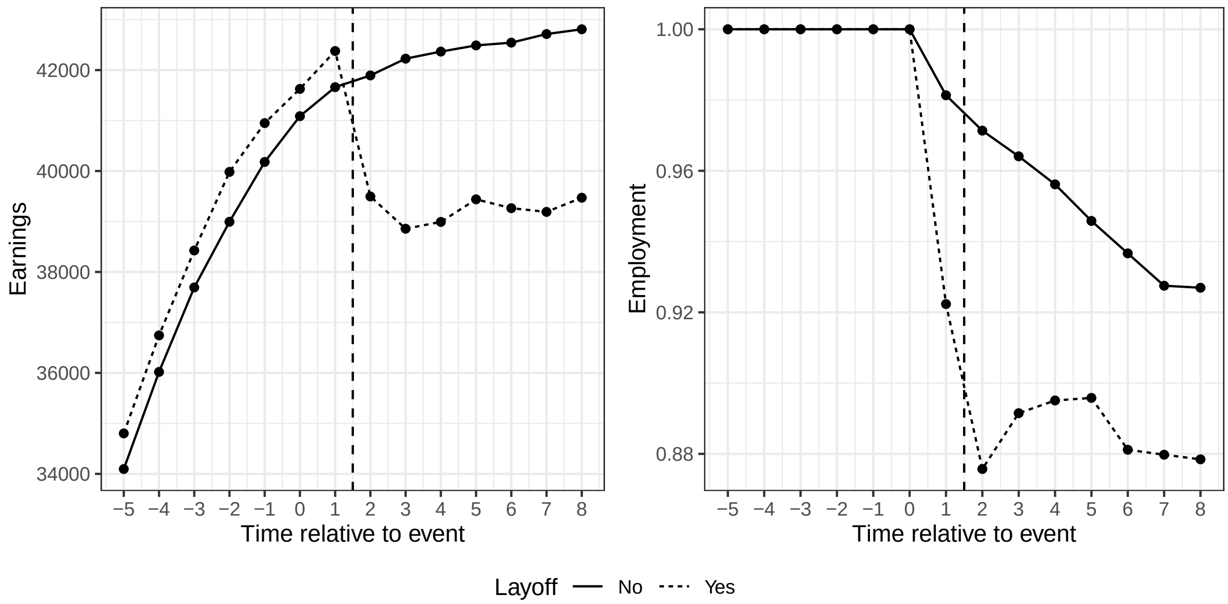

Figure B-4: Baseline Raw Means for the Mass-Layoff Design.

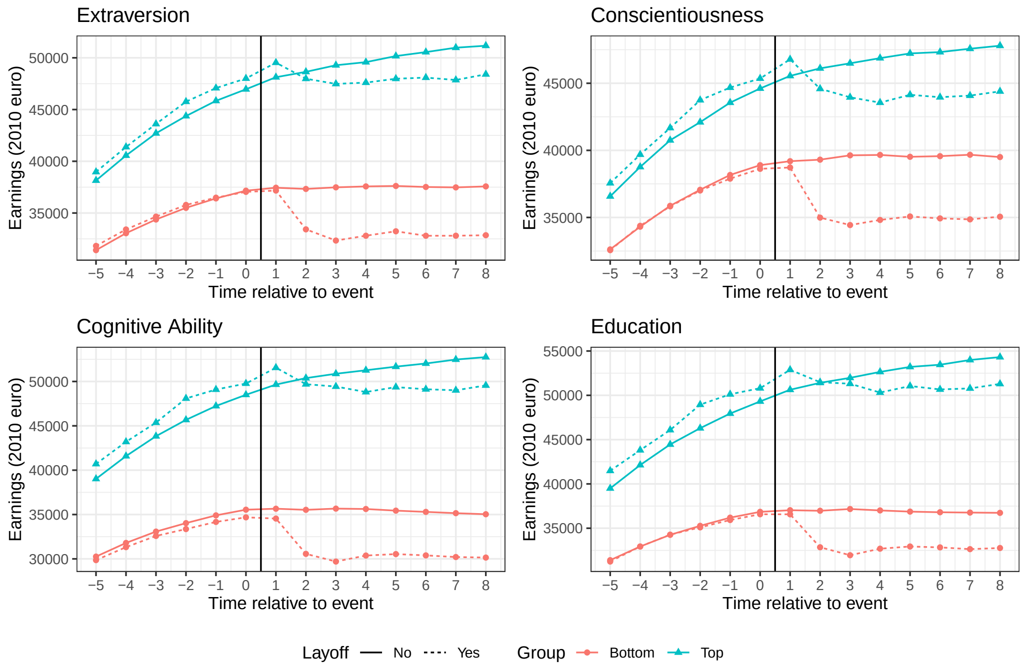

(a) Earnings.

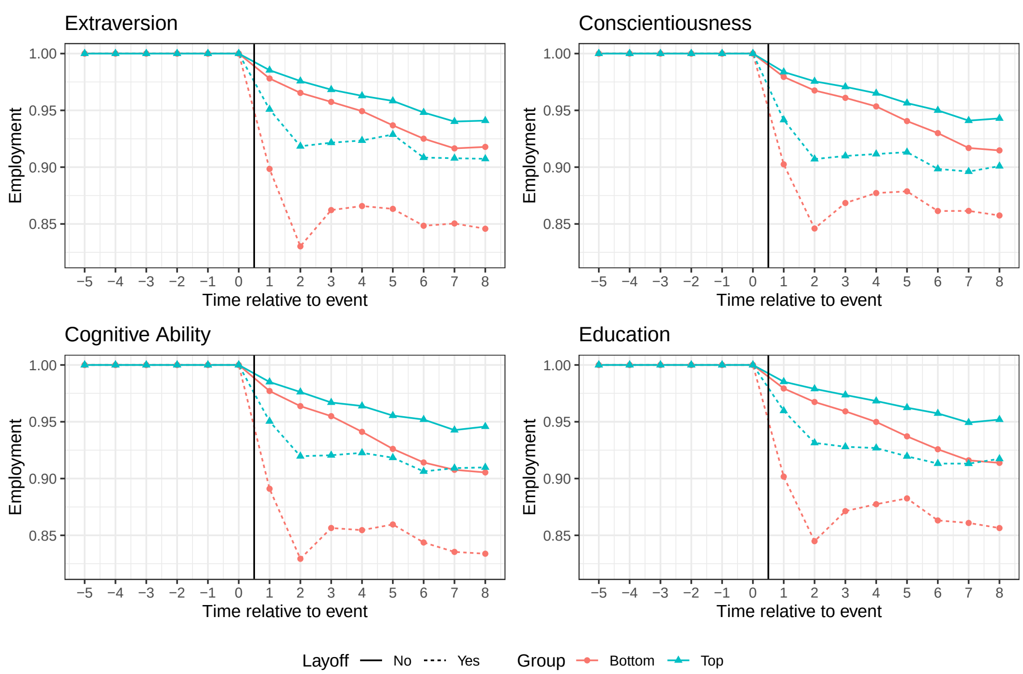

(b) Employment.

Figure B-5: Multidimensional Raw Quartile Means.

Notes: The figure shows the labor market outcomes for workers who experience a mass layoff or plant closure event and those who do not. The dashed lines correspond to workers who experienced an event between periods 1 and 2, while the solid lines correspond to the matched control group with no event. Within those groups, the blue line corresponds to workers in the top quartile of the indicated trait, while the red line corresponds to the bottom quartile.

Table B.1: Cross-Correlations: Raw Traits.

Correlations (Full Sample)

	IQ Visuospatial	IQ Math	IQ Verbal	Social	Activity	Masculinity	Deliberation	Dutifullness	Achievement Aim	Leadership	Confidence	HS Verbal	HS English	HS Math Adv.	HS Math Basic	HS Average	School GPA	Earnings at 36
IQ Visuospatial	1	0.65	0.6	0.17	0.16	0.03	0.14	0.23	0.29	0.26	0.29	0.18	0.24	0.26	0.26	0.27	0.44	0.25
IQ Math	0.65	1	0.71	0.18	0.18	0.04	0.15	0.25	0.34	0.28	0.31	0.3	0.3	0.38	0.36	0.41	0.6	0.31
IQ Verbal	0.6	0.71	1	0.2	0.18	0.01	0.16	0.28	0.34	0.3	0.32	0.35	0.35	0.27	0.22	0.41	0.58	0.27
Social	0.17	0.18	0.2	1	0.59	0.12	0.19	0.44	0.47	0.75	0.7	−0.02	−0.02	−0.06	−0.03	−0.03	0.11	0.19
Activity	0.16	0.18	0.18	0.59	1	0.28	0.38	0.49	0.6	0.66	0.61	−0.04	−0.09	−0.03	∘	−0.04	0.11	0.19
Masculinity	0.03	0.04	0.01	0.12	0.28	1	0.1	0.03	0.1	0.09	0.21	−0.15	−0.17	−0.08	−0.05	−0.16	−0.14	0.02
Deliberation	0.14	0.15	0.16	0.19	0.38	0.1	1	0.64	0.38	0.31	0.41	0.15	0.04	0.14	0.08	0.16	0.27	0.15
Dutifullness	0.23	0.25	0.28	0.44	0.49	0.03	0.64	1	0.54	0.56	0.51	0.15	0.07	0.11	0.06	0.16	0.36	0.2
Achievement Aim	0.29	0.34	0.34	0.47	0.6	0.1	0.38	0.54	1	0.69	0.48	0.14	0.09	0.14	0.09	0.18	0.37	0.25
Leadership	0.26	0.28	0.3	0.75	0.66	0.09	0.31	0.56	0.69	1	0.64	0.08	0.06	0.04	0.04	0.09	0.26	0.25
Confidence	0.29	0.31	0.32	0.7	0.61	0.21	0.41	0.51	0.48	0.64	1	0.07	0.05	0.08	0.05	0.1	0.26	0.22
HS Verbal	0.18	0.3	0.35	−0.02	−0.04	−0.15	0.15	0.15	0.14	0.08	0.07	1	0.51	0.43	0.33	0.78	0.54	0.18
HS English	0.24	0.3	0.35	−0.02	−0.09	−0.17	0.04	0.07	0.09	0.06	0.05	0.51	1	0.42	0.3	0.79	0.45	0.18
HS Math Adv.	0.26	0.38	0.27	−0.06	−0.03	−0.08	0.14	0.11	0.14	0.04	0.08	0.43	0.42	1	0.47	0.75	0.49	0.23
HS Math Basic	0.26	0.36	0.22	−0.03	∘	−0.05	0.08	0.06	0.09	0.04	0.05	0.33	0.3	0.47	1	0.65	0.4	0.14
HS Average	0.27	0.41	0.41	−0.03	−0.04	−0.16	0.16	0.16	0.18	0.09	0.1	0.78	0.79	0.75	0.65	1	0.66	0.26
School GPA	0.44	0.6	0.58	0.11	0.11	−0.14	0.27	0.36	0.37	0.26	0.26	0.54	0.45	0.49	0.4	0.66	1	0.35
Earnings at 36	0.25	0.31	0.27	0.19	0.19	0.02	0.15	0.2	0.25	0.25	0.22	0.18	0.18	0.23	0.14	0.26	0.35	1

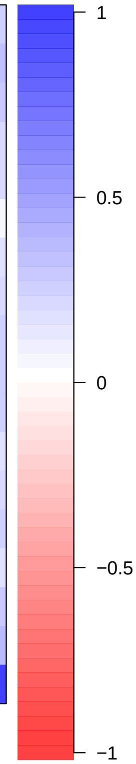

Notes: Data sources described in Section 2.2 and Appendix B.2.

Table B.2: Factor Loadings.

Factor Loadings

Variable	MR1	MR2	MR3	h^2	u^2	com
Sociability	**0.91**	-0.05	-0.14	0.71	0.29	1.06
Leadership	**0.87**	0.05	-0.01	0.79	0.21	1.01
Activity-Energy	**0.74**	-0.07	0.14	0.62	0.38	1.10
Confidence	**0.69**	0.10	0.12	0.62	0.38	1.10
Achivement Aim	**0.54**	0.17	0.19	0.53	0.47	1.46
Masculinity	0.20	-0.05	0.02	0.04	0.96	1.14
Deliberation	-0.03	-0.02	**0.94**	0.86	0.14	1.00
Dutifulness	0.34	0.08	**0.53**	0.60	0.40	1.76
Arithmetic	-0.02	**0.88**	-0.02	0.76	0.24	1.00
Verbal	0.01	**0.81**	0.00	0.66	0.34	1.00
Visuospatial	0.00	**0.75**	-0.01	0.55	0.45	1.00
SS loadings	3.24	2.09	1.41			
MR1	1.00	0.35	0.43			
MR2	0.35	1.00	0.24			
MR3	0.43	0.24	1.00			

Notes: Oblique rotation is used to obtain loadings. MR1 (MinRes solution) is labeled Extraversion, MR2 is labeled Cognitive Ability and MR3 is labeled Conscientiousness.

Table B.3: Balance Table: Workers.

Variable	Treat. Mean	Control Mean	Mean Dif.	t stat.	Treat. N	Control N
Earnings	41,600	41,100	540.6	-4.4	17,581	86,373
Age	39.2	39.2	-0.02	0.9	17,581	86,373
Tenure	9.3	9.5	-0.2	4.4	17,581	86,373
Plant size	296.8	276.4	20.4	-6.1	17,581	86,373
Years of Education	12.7	12.7	0.03	-1.6	17,581	86,373
College Educated	0.4	0.3	0.01	-3.8	17,581	86,373
Extraversion	-0.02	-0.1	0.04	-5.1	17,581	86,373
Conscientiousness	0.1	0.03	0.03	-3.2	17,581	86,373
Cognitive ability	0.1	0.002	0.1	-7.1	17,581	86,373

Notes: Each column reports a summary number of the indicated variable across all establishments in period zero. The first three report the total number of closures and mass layoffs occurring in the treatment group. Plant Size is the average number of employees in period zero.

Table B.4: Balance Table: Firms.

Group	Events	Closures	Mass Layoffs	Aveg. Emp. Change	Plant Size
Treatment	3,535	1,118	2,417	0.50	68.3
Control	0	0	0	1.04	51.7

Notes: Columns indicate the means of the row variables in the treatment and control groups in period zero. The mean difference between the treatment and the control groups and its associated t-statistic is also shown. Firm's Employment Change is the average firm growth from period 0 to period 1. Plant Size is the average number of employees in period zero. Tenure is the number of consecutive years employed in the period zero establishment.

Table B.5: Cross-Sectional Estimates: Matched Sample, Pre-Period.

Dependent Variable:	log(Earnings)				
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Extraversion	0.086				
	(0.003)				
Conscientiousness		0.048			
		(0.003)			
Cognitive Ability			0.119		
			(0.004)		
Years of Education				0.066	
				(0.002)	
Age					0.026
					(0.001)
Outcome mean	10.6	10.6	10.6	10.6	10.6
Fixed-effects					
Birth Year (13)	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	103,954	103,954	103,954	103,954	103,954
R^2	0.08060	0.03606	0.13453	0.19735	0.05829
Within R^2	0.06584	0.02058	0.12063	0.18446	0.04317

Notes: Each column reports the OLS regressions results from Equation 2.1 with log earnings as the outcome in the matched sample in the pre period. The unit of observation is the person. Extraversion, conscientiousness, and cognitive ability are constructed using exploratory factor analysis and normalized to have mean 0 and standard deviation 1 within cohorts. Years of education is constructed by mapping the highest degree at age 35 to its official length (e.g., a high-school degree equals 12 years of education). Heteroskedasticity-robust standard-errors are in parentheses.

B.2 Supplementary Data Description

B.2.1 The Finnish Defence Forces (FDF) Test Data

B.2.1.1 Background

Military conscription in Finland is universal and grants relatively few exceptions. The available data cover 80% of Finnish men born between 1962 and 1979 (n = 489,252). Finnish men are drafted in the year they turn 18 and most start their service at age 19 or 20. Military service lasts for 6–12 months, and most conscripts do not continue service at the military.

FDF uses psychological tests as of the criteria to assess conscripts' suitability for non-commissioned officer training. FDF conducted psychological tests on all conscripts since 1955. Between 1955 and 1982, FDF used one test that measured cognitive skills: logical, mathematical and verbal skills. From 1982, the FDF has used two tests: a cognitive and a personality test. The content of each test is described in the sections below.

The test data have been described in Jokela et al. (2017) and validated in FDF's internal reports summarized in Nyman (2007).

B.2.1.2 Administration of the Tests

The cognitive ability test and the personality test are typically taken in the second week of military service in a 2-h paper-and-pencil format in standardized group-administered conditions. The personality test contains 218 statements with a response scale of yes/no. The cognitive test contains 120 multiple-choice questions. The test questionnaires have been unchanged from 1982 to 2000 (the data available to this study), and the scores are designed to be comparable across cohorts. The main change in the test administration during the timeline of this study is that between 1995 and 2000, the personality test was administered already at the conscription, on average 18 months before entering the FDF service. The administration of the cognitive test has been unchanged 1982–2000.

B.2.1.3 Selection Concerns

The data are subject to two selection concerns. The first concern is selection into military service: Only those that enter the FDF service take the tests. It is possible to be exempted from the military service due to severe health conditions, most often related to mental health problems, or due to religious or ethical convictions. For the analysis, this means that the sample is generally more representative of men with relatively higher labor-market prospects. Over the timeline of this study, selection into military service has been stable (Jokela et al., 2017). The second concern is the selective test performance. The military uses the test results for selecting conscripts to officer training. To some extent, this feature is likely to induce higher performance from those that would like to be selected and lower performance that would like to avoid it. For personality data, the concern is alleviated by the fact that the scoring rules are not revealed to the conscripts. For cognitive data, test performance may reflect, to some extent, motivation-related factors, as is the case for most cognitive tests. Finally, the data excludes The Finnish Defense Forces personnel as well as Finnish Border Guard soldiers.

B.2.1.4 The Cognitive Ability Test

The cognitive ability test has three subtests: visuospatial, arithmetic and verbal reasoning. The FDF cognitive ability test is similar to the The Armed Services Vocational Aptitude Battery (ASVAB), administered by the United States Military Entrance Processing Command. Each subtest has 40 multiple-choice questions. FDF reports test-retest reliabilities of the subtests between 0.76 and 0.88 (Nyman, 2007). The descriptions of tests are based on Nyman (2007) and Jokela et al. (2017):

1. *The visuospatial subtest* is similar to Raven's Progressive Matrices Raven et al. 2000. The test shows a set of matrices, each with one removed part, and the participant choose a figure that completes the matrix.

2. *The arithmetic subtest* contains different tasks: computing arithmetic operations, completing a series of numbers that follow a pattern, solving short verbal problems, and noticing similarities in relationships between numbers.

3. *The verbal subtest* requires choosing synonyms or antonyms, selecting a word that belongs to the same category as the given pair of words, choosing which word on a list does not belong in the group, and detecting similar relationships between two pairs of words (Jokela et al., 2017).

B.2.1.5 The Personality Test

The personality test aims to measure 8 personality traits. The test is similar to and partly based on the Minnesota Multiphasic Personality Inventory (MMPI). It contains 218 statements with a yes/no response scale—between 18 and 33 items for each personality trait. The test score for each personality trait is the sum of the binary answers aligned with the trait (for example, in reverse-coded statements, cases where the task-taker disagrees). The data available to this study contain these sums of scores. FDF reports that internal reliability varies between 0.6 and 0.9 by trait and that the average Cronbach alpha is 0.75 (Nyman, 2007).

The 8 personality traits measured in the test are, as described by Jokela et al. (2017):

1. *Sociability*: the person's level of gregariousness and preference for socializing with others (33 items; e.g., whether the person likes to host parties and not withdraw from social events).

2. *Activity–energy*: how much the person exerts physical effort in everyday activities and how quickly the person prefers to execute activities (28 items; e.g., whether the person tends to work fast and vigorously and prefers fast-paced work).

3. *Self-confidence*: the person's self-esteem and beliefs about his abilities (32 items; e.g., whether the person feels to be as good and able as others and can meet other people's expectations).

4. *Leadership motivation*: how much the person prefers to take charge in groups and influence other people; it includes 30 items.

5. *Achievement motivation*: how strongly the person wants to perform well and achieve important life goals (24 items; e.g., whether the person is prepared to make personal sacrifices to achieve success).

6. *Dutifulness*: how closely the person follows social norms and considers them to be important (18 items; e.g., whether the person would return money if given back too much change at a store).

7. *Deliberation*: how much the person prefers to think ahead and plan things before acting (26 items; e.g., whether the person prefers to spend money carefully).

8. *Masculinity*: the person's occupational and recreational interests that are traditionally considered as masculine (27 items; e.g., whether the person would like to work as a construction manager).

Dutifulness, deliberation, achievement striving are all related to the higher order personality factor conscientiousness.

The FDF personality test also includes questions about mental health and questions targeted at evaluating the answers' validity. The mental health part has four mental health sub-scales from the Minnesota Multiphasic Personality Inventory (MMPI) as described by Psych Central (retrieved 2020):

1. *Psychopathic deviate*: General social maladjustment and the absence of pleasant experiences. Associated with narcissism, externalization of blame, exploitativeness, and hostility.

2. *Psychasthenia*: Person's inability to resist specific actions or thoughts, regardless of their maladaptive nature. "Psychasthenia" is an old term used to describe a phenomenon that is currently called obsessive-compulsive disorder (OCD).

3. *Schizophrenia*: Bizarre thoughts, peculiar perceptions, social alienation, poor familial relationships, difficulties in concentration and impulse control, lack of deep interests, disturbing question of self-interest and self-worth, and sexual difficulties.

4. *Hypochondriasis*: Wide variety of vague and non-specific complaints about bodily functioning. Complaints tend to focus on back and abdomen, and they persist in the face of negative medical tests.

The validity part has three sub-scales as:

1. *L-scale*: Attempts to give an overly favorable impression of one's conduct; persons' test-taking attitude and approach to the test: intended to identify people who deliberately try to avoid answering the test honestly and in a frank manner.

2. *K-scale*: Persons' test-taking attitude and approach to the test: designed to identify psychopathology in people who otherwise would have profiles within the normal range. A subtle measure: high scores combined with prior information on psychological problems are interpreted as a signal of defensiveness. High-scores without previous psychological problems tend to be observed with confident individuals.

3. *F-scale*: Attempts to give unusual, for example, random or contradictory answers; persons' test-taking attitude and approach to the test: intended to detect unusual or atypical ways of answering the test items.

B.2.1.6 Exploratory Factor Analysis

The raw data provide test scores for 8 personality dimensions, 3 cognitive-skill dimensions, 4 psychopathological dimensions, and 3 test validity measures. We first consider only the personality and cognitive-skill test scores. The cross-correlation matrix in Table B.1 shows that both personality and cognitive measures are correlated within their domains. Within personality scores, the cross-correlation matrix suggests that the traits with labels related to extraversion (sociability, activity, confidence, and leadership) and conscientiousness (deliberation and dutifulness) have relatively higher correlations within their subdomains.[1] Achievement aim is traditionally associated with conscientiousness but in the FDF test, it has relatively high correlations also with the extraversive traits. Masculinity has low correlations with other personality traits and cognitive measures.

We also expanded the set of variables by including the psychopathological dimensions and test validity measures, each in turn. In the expanded four-factor model, the psychopathological dimensions load together into single factor, separate from cognitive, extraversive, and conscientiousness-related factors. However, self-confidence now loads into the psychopathological factor with a negative loading, and we note that the psychopathological factor is relatively strongly correlated ($\rho = .6$) with the extraversive factor. We infer that the psychopathological factor captures many aspects of the extraversion-related factor. This observation is also supported by regression evidence, where including both in a regression typically leads to a coefficient of close to zero for the other. We decide not to include the psychopathological measures in our main factorization because (1) it contains limited variations, (2) the evidence indicates that it is uncertain whether the measure is sufficiently separate from the extraversion-related factor, and (3) we want focus on the distinction between interpersonal and intrapersonal skill.

[1]Extraversion and conscientiousness are elements of the Big Five and five-factor personality models. Extraversion is also one of the three personality dimensions in Eysenck's dimensions.

Appendix C

Appendix to Chapter 3

C.1 Data

C.1.1 The Finnish Defence Forces (FDF) Test Data

Background Military conscription in Finland is universal and grants relatively few exceptions. The available data cover 80% of Finnish men born between 1962 and 1979 (n = 489,252). Finnish men are drafted in the year they turn 18 and most start their service at age 19 or 20. Military service lasts for 6–12 months, and most conscripts do not continue service at the military.

FDF uses psychological tests as of the criteria to assess conscripts' suitability for non-commissioned officer training. FDF conducted psychological tests on all conscripts since 1955. Between 1955 and 1982, FDF used one test that measured cognitive skills: logical, mathematical and verbal skills. From 1982, the FDF has used two tests: a cognitive and a personality test. The content of each test is described in the sections below.

The test data have been described in Jokela et al. (2017) and validated in FDF's internal reports summarized in Nyman (2007).

Administration of the Tests The cognitive ability test and the personality test are typically taken in the second week of military service in a 2-h paper-and-pencil format in standardized group-administered conditions. The personality test contains 218 statements with a response scale of yes/no. The cognitive test contains 120 multiple-choice questions. The test questionnaires have been unchanged from 1982 to 2000 (the data available to this study), and the scores are designed to be comparable across cohorts. The main change in the test administration during the timeline of this study is that between 1995 and 2000, the personality test was administered already at the conscription, on average 18 months before entering the FDF service. The administration of the cognitive test has been unchanged 1982–2000.

The Cognitive Ability Test The cognitive ability test has three subtests: visuospatial, arithmetic and verbal reasoning. The FDF cognitive ability test is similar to the The Armed Services

Vocational Aptitude Battery (ASVAB), administered by the United States Military Entrance Processing Command. Each subtest has 40 multiple-choice questions. FDF reports test–retest reliabilities of the subtests between 0.76 and 0.88 (Nyman, 2007). The descriptions of tests are based on Nyman (2007) and Jokela et al. (2017):

1. *The visuospatial subtest* is similar to Raven's Progressive Matrices (Raven and Court, 1938). The test shows a set of matrices, each with one removed part, and the participant choose a figure that completes the matrix.

2. *The arithmetic subtest* contains different tasks: computing arithmetic operations, completing a series of numbers that follow a pattern, solving short verbal problems, and noticing similarities in relationships between numbers.

3. *The verbal subtest* requires choosing synonyms or antonyms, selecting a word that belongs to the same category as the given pair of words, choosing which word on a list does not belong in the group, and detecting similar relationships between two pairs of words (Jokela et al., 2017).

The Personality Test The personality test aims to measure 8 personality traits. The test is similar to and partly based on the Minnesota Multiphasic Personality Inventory (MMPI). It contains 218 statements with a yes/no response scale—between 18 and 33 items for each personality trait. The test score for each personality trait is the sum of the binary answers aligned with the trait (for example, in reverse-coded statements, cases where the task-taker disagrees). The data available to this study contain these sums of scores. FDF reports that internal reliability varies between 0.6 and 0.9 by trait and that the average Cronbach alpha is 0.75 (Nyman, 2007).

The 8 personality traits measured in the test are, as described by Jokela et al. (2017):

1. *Sociability*: the person's level of gregariousness and preference for socializing with others (33 items; e.g., whether the person likes to host parties and not withdraw from social events).

2. *Activity–energy*: how much the person exerts physical effort in everyday activities and how quickly the person prefers to execute activities (28 items; e.g., whether the person tends to work fast and vigorously and prefers fast-paced work).

3. *Masculinity*: the person's occupational and recreational interests that are traditionally considered as masculine (27 items; e.g., whether the person would like to work as a construction manager).

4. *Dutifulness*: how closely the person follows social norms and considers them to be important (18 items; e.g., whether the person would return money if given back too much change at a store).

5. *Deliberation*: how much the person prefers to think ahead and plan things before acting (26 items; e.g., whether the person prefers to spend money carefully).

6. *Achievement motivation*: how strongly the person wants to perform well and achieve important life goals (24 items; e.g., whether the person is prepared to make personal sacrifices to achieve success).

7. *Leadership motivation*: how much the person prefers to take charge in groups and influence other people; it includes 30 items.

8. *Self-confidence*: the person's self-esteem and beliefs about his abilities (32 items; e.g., whether the person feels to be as good and able as others and can meet other people's expectations).

Dutifulness, deliberation, achievement striving are all related to the higher order personality factor conscientiousness.

The FDF personality test also includes questions about mental health and questions targeted at evaluating the answers' validity. The mental health part has four mental health sub-scales from the Minnesota Multiphasic Personality Inventory (MMPI), hypochondriasis, psychopathic deviate, psychasthenia, and schizophrenia. The validity part has three sub-scales: lie (attempts to give an overly favorable impression of one's conduct), fix (attempts to give an overly unfavorable impression of one's conduct), and validity (attempts to give unusual, for example, random or contradictory answers).

Selection Concerns The data are subject to two selection concerns. The first concern is selection into military service: Only those that enter the FDF service take the tests. It is possible to be exempted from the military service due to severe health conditions, most often related to mental health problems, or due to religious or ethical convictions. For the analysis, this means that the sample is generally more representative of men with relatively higher labor-market prospects. Over the timeline of this study, selection into military service has been stable (Jokela et al., 2017). The second concern is the selective test performance. The military uses the test results for selecting conscripts to officer training. To some extent, this feature is likely to induce higher performance from those that would like to be selected and lower performance that would like to avoid it. For personality data, the concern is alleviated by the fact that the scoring rules are not revealed to the conscripts. For cognitive data, test performance may reflect, to some extent, motivation-related factors, as is the case for most cognitive tests. Finally, the data excludes The Finnish Defense Forces personnel as well as Finnish Border Guard soldiers.

C.1.2 Anchoring High-School Test Data

In high school, students can select between two tracks of mathematics; basic and advanced. The exit exams are different for both tracks and a small fraction opts out from both. Our aim is to construct a single measure of mathematics test scores that is commensurable across the three options. We do this by regressing:

$$\begin{aligned}\text{MilitaryMathScore}_{it} = {} & \delta_1 D_i^{\text{BasicMath}} + \delta_2 D_i^{\text{AdvancedMath}} \\ & + \delta_3 D_i^{\text{BasicMath}} \text{BasicMathScore}_i \\ & + \delta_4 D_i^{\text{AdvancedMath}} \text{AdvancedMathScore}_i + \delta_t \end{aligned} \tag{C.1}$$

where D indicates that person i has participated in the exam. The indicator is interacted with the normalized test score. For those who did not participate, number -1 is imputed for the test score (the scalar used here does not matter for the estimation). Finally, δ_t is a fixed effect for the test-taking year.

The left-hand side variable is the military arithmetic test score. We use the fact that this standardized test is administered to everyone in our data. The military test is low stakes, relatively easy, and only moderately correlated with the high-school test scores (less than 0.4 with either track). While it does not share the same patterns as our main results (results not shown), it is a reasonable tool for this purpose.

Table C.1 shows the estimation results. The marginal weights for better test scores are similar in both tracks. Both predict around 0.26 standard deviation increase in the military test for each standard deviation increase in the high school score. The differences arise from a mean shift in the military arithmetic test. The mean performance of students taking the advanced mathematics track is almost 0.5 standard deviations higher than the mean performance of students taking the basic track ($\delta_2 - \delta_1$). The 'math' variable in all results except for Table 3.1 is the weighted average of the right hand side variables, where the weights are given by the δ values. The cohort fixed effects are not included.

Table C.1: Math Anchoring Regressions.

	Outcome: Military Arithmetic Test
δ_1	0.334
	(0.006)
δ_2	0.810
	(0.005)
δ_3	0.255
	(0.003)
δ_4	0.270
	(0.002)
Num. obs.	165934
Adj. R^2 (full model)	0.269
Adj. R^2 (proj model)	0.255

Notes: Robust standard errors are in parentheses.

C.2 Proofs

Optimal time allocation s^*

Proof. Using equation 3.5:

$$\begin{aligned}\frac{\partial U(s;N,J)}{\partial s} = -a(N,J) + b(N,J) + \frac{\partial V(s;N,J)}{\partial s} &= 0\\ \frac{\partial C(s;N,J)}{\partial s} &= a(N,J) - b(N,J)\\ s^* &= g_s(a(N,J) - b(N,J);N,J)\end{aligned}$$

■

Comparative static: Optimal response of s to a change in J

Proof. Differentiate with respect to J from both sides of equation 3.5:

$$\begin{aligned}\frac{\partial^2 V(.)}{\partial s^2}\frac{\partial s^*}{\partial J} + \frac{\partial^2 V(.)}{\partial s^* \partial J} &= a^J(N,J) - b^J(N,J)\\ -\frac{\partial^2 V(.)}{\partial s^2}\frac{\partial s^*}{\partial J} &= b^J(N,J) + \frac{\partial^2 V(.)}{\partial s^* \partial J} - a^J(N,J)\\ \frac{\partial s^*}{\partial J} &= -\left[\frac{\partial^2 V(.)}{\partial s^2}\right]^{-1}\left[b^J(N,J) + \frac{\partial^2 V(.)}{\partial s^* \partial J} - a^J(N,J)\right]\end{aligned}$$

■

Comparative static: Marginal returns to a change in J

Proof. Earnings are given by

$$\begin{aligned}Y &= r_H H(1-s;N,J) + r_S S(s;N,J) + r_N N + r_J J\\ &= r_H a(N,J)(1 - s^*(N,J)) + r_S b(N,J) s^*(N,J) + r_N N + r_J J\\ &= r_H a(N) - r_H a(N) s^*(N,J) + r_S b(J) s^*(N,J) + r_N N + r_J J\end{aligned}$$

Differentiate with respect to J:

$$\begin{aligned}\frac{\partial Y}{\partial J} =& - r_H a(N)\frac{\partial s^*(N,J)}{\partial J}\\ &+ r_S\left[b'(J)s^*(N,J) + b(J)\frac{\partial s^*(N,J)}{\partial J}\right] + r_J\\ &= \underbrace{r_S b'(J) s^*(N,J) + r_J}_{\text{direct effect}} + \underbrace{(r_S b(J) - r_H a(N))}_{\text{net earnings change}}\underbrace{\frac{\partial s^*(N,J)}{\partial J}}_{\text{change in } s}\end{aligned}$$

■

C.3 Robustness

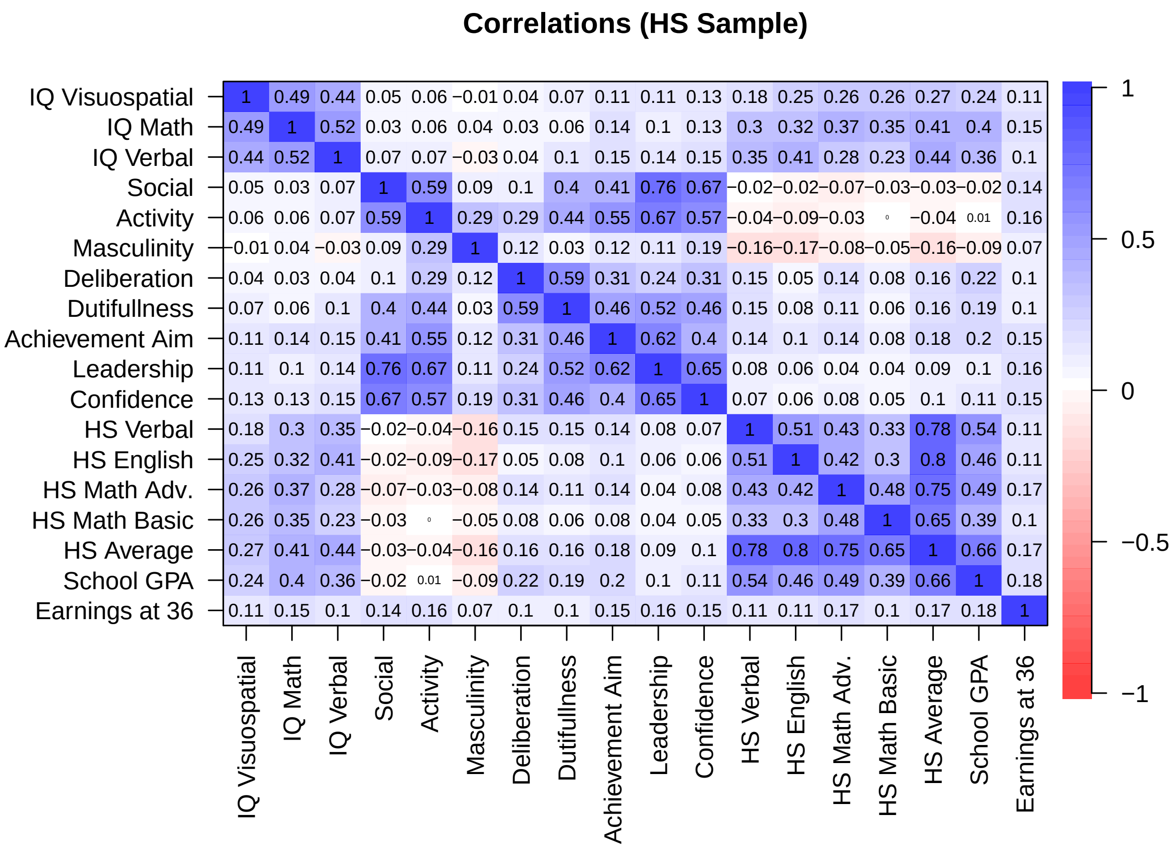

Figure C-1: Cross-correlations.

Notes: Each number is a pairwise correlation coefficient with person as the unit of observation. All variables are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are recorded by the tax authorities and measured by averaging total labor and entrepreneurial income earned at age 35-38. The data includes only persons for which we have high-school data.

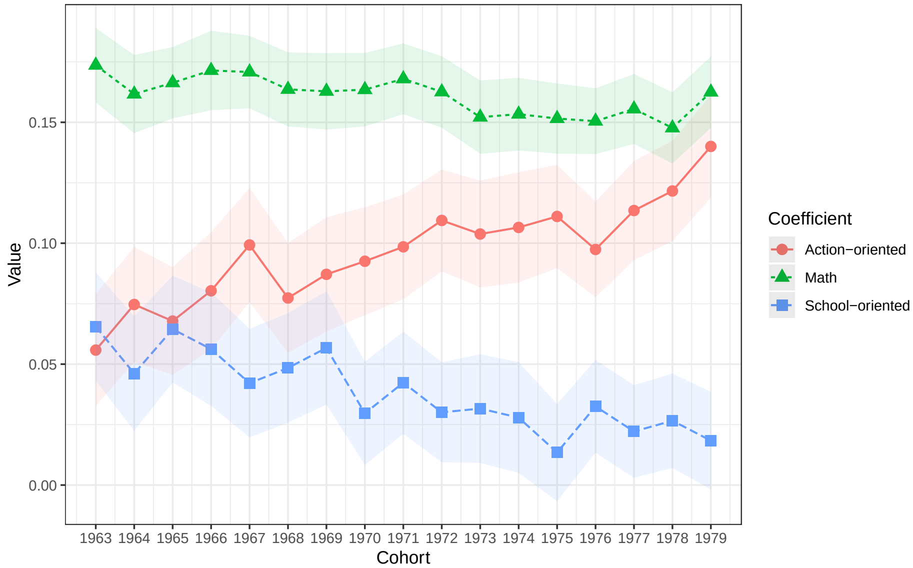

Figure C-2: Trends in Trait Premia with Mathematics.

Notes: Each point in the figure corresponds to a regression coefficient from estimating Equation 3.13 separately for each cohort, with log earnings as the outcome and person as the unit of observation. The right-hand-side variables include only the action-oriented and school-oriented traits. The action-oriented trait is a composite of Sociability, Activity, and Masculinity. The school-oriented trait is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. All covariates are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are recorded by the tax authorities and measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported as the shaded area.

Table C.2: Returns to Skills.

	Dependent variable: log earnings		
	(1)	(2)	(2)
Sociability	0.019		0.065
	(0.003)		(0.003)
Activity	0.023		0.047
	(0.003)		(0.003)
Masculinity	0.023		0.035
	(0.002)		(0.002)
Deliberation	0.041		0.021
	(0.002)		(0.002)
Dutifulness	−0.032		−0.037
	(0.003)		(0.003)
Achievement aim	0.050		0.017
	(0.002)		(0.002)
Confidence	0.037		0.000
	(0.003)		(0.003)
Leadership	0.038		0.021
	(0.003)		(0.003)
Math		0.138	0.135
		(0.002)	(0.002)
Verbal		−0.006	0.005
		(0.002)	(0.002)
Electives		0.056	0.052
		(0.002)	(0.002)
Outcome mean	10.520	10.520	10.520
Cohort FE	yes	yes	yes
Adj. R^2	0.050	0.068	0.099
Observations	157743	157605	156843

Notes: Each column reports the OLS regression results from Equation 3.13, with log earnings as the outcome. The unit of observation is the person. 'Action-oriented' is a composite of Sociability, Activity, and Masculinity. 'School-oriented' is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. Test scores and traits are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported in parentheses.

Table C.3: Returns to Skills: Levels.

	Dependent variable: Earnings (2010 euros)			
	(1)	(2)	(3)	(4)
Action-oriented	1051.077		2939.034	3678.742
	(77.473)		(76.447)	(78.739)
School-oriented	4816.418		2367.299	1461.602
	(77.647)		(76.984)	(79.139)
Math		6829.824	6480.433	5102.159
		(55.925)	(56.176)	(66.391)
IQ				52.229
				(58.347)
Verbal				690.042
				(63.882)
Electives				2539.078
				(73.595)
Outcome mean	44325	44290	44328	44350
Cohort FE	yes	yes	yes	yes
Adj. R^2	0.076	0.105	0.150	0.160
Observations	157743	157891	157129	156843

Notes: Each column reports the OLS regression results from Equation 3.13, with earnings in 2010 euros as the outcome. The unit of observation is the person. 'Action-oriented' is a composite of Sociability, Activity, and Masculinity. 'School-oriented' is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. Test scores and traits are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported in parentheses.

Table C.4: Returns to Skills: Truncated Sample.

	Dependent variable: log(Earnings)			
	(1)	(2)	(3)	(4)
Action-oriented	0.034		0.078	0.092
	(0.002)		(0.002)	(0.002)
School-oriented	0.101		0.044	0.026
	(0.002)		(0.002)	(0.002)
Math		0.156	0.150	0.122
		(0.001)	(0.001)	(0.002)
IQ				0.007
				(0.002)
Verbal				0.009
				(0.002)
Electives				0.051
				(0.002)
Outcome mean	10.570	10.570	10.570	10.570
Cohort FE	yes	yes	yes	yes
Adj. R^2	0.065	0.090	0.127	0.134
Observations	155704	155840	155097	154822

Notes: Each column reports the OLS regression results from Equation 3.13, with log earnings as the outcome. The model is estimated with truncated data using log earnings > 8 as the threshold. 'Action-oriented' is a composite of Sociability, Activity, and Masculinity. 'School-oriented' is a composite of Deliberation, Dutifulness, Achievement aim, Confidence, and Leadership. Test scores and traits are normalized to have mean 0 and standard deviation 1 within cohorts. Earnings are measured by averaging total labor and entrepreneurial income earned at age 35-38. Robust standard errors are reported in parentheses.

Appendix D

Appendix to Chapter 4

Table D.1: Alternative IV: 2SLS Estimates of Manufacturing Decline Effects on High-School Dropout Rate 1991–2011 and College Mobility 1999–2011.

Alternative 2SLS Estimates	(1)	(2)	(3)
	A. High-School Dropout Rate		
Commuting zone manufacturing decline	-.270***	-.166***	-.441*
	(.038)	(.053)	(.235)
Other baseline controls	–	Yes	Yes
Baseline manufacturing emp. share	–	–	Yes
	B. College Mobility		
Commuting zone manufacturing decline	.562***	.438***	.587
	(.125)	(.151)	(.393)
Other baseline controls	–	Yes	Yes
Baseline manufacturing emp. share	–	–	Yes
	C. 2SLS First Stage Estimates†		
Commuting zone import exposure	-2.29***	-1.87***	-.63***
	(.23)	(.13)	(.016)
F-statistic	155.0	48.4	71.7
Adjusted R^2	0.22	0.27	0.39

Notes: Alternative IV specification. In Panel A, each column reports results from stacking changes in commuting zone high-school dropout rates and declines in manufacturing-to-total employment ratios over the periods 1991–99 and 1999–2011. The dependent variable is the annual change in the high-school dropout rate (N = 1,444 = 722 commuting zones × 2 periods). High-school dropout rate is computed from the US Census for 1990 and 2000, and from the ACS for 2011 as a five-year average. In Panel B, each column reports results from regressing changes in commuting zone measures of absolute college mobility on declines in manufacturing-to-total employment ratios over the period 1999–2011. The dependent variable is the annual change in college mobility between cohorts born in 1984 and 1993 (N = 616 commuting zones). College mobility is CZ-level average of college attendance of children with parents at the 25th percentile in the national distribution. The college mobility measure comes from Chetty et al. (2014) and is based on the US tax records. In Panels A and B, manufacturing decline is instrumented with an alternative measure of the commuting zone import exposure, constructed from Chinese imports to eight other high-income countries, excluding the US, as in Autor et al. (2013) and detailed in the text. The commuting zone baseline manufacturing controls are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 period. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. Other baseline controls include population counts, employment-to-population ratios, and region controls for nine regional census divisions. All models in Panel A include a time trend. Standard errors are clustered by commuting zone.
† For manufacturing share over 1991–2011.
$*p < 0.10$
$**p < 0.05$
$***p < 0.01$

Table D.2: The Reduced Form: OLS and 2SLS Estimates of Trade Exposure Effects on High-School Dropout Rate 1991–2011 and College Mobility 1999–2011.

Reduced Form Estimates	OLS		Combined 2SLS	
	(1)	(2)	(3)	(4)
	A. High-School Dropout Rate			
Commuting zone import exposure	-.357***	-.338***	-.543***	-.656**
	(.082)	(.106)	(.171)	(.295)
Baseline manufacturing emp. share	–	Yes	–	Yes
Other baseline controls	Yes	Yes	Yes	Yes
	B. College Mobility			
Commuting zone import exposure	.674*	.016	1.41***	1.23
	(.392)	(.474)	(.493)	(.815)
Baseline manufacturing emp. share	–	Yes	–	Yes
Other baseline controls	Yes	Yes	Yes	Yes

Notes: Reduced form regression. In Panel A, each column reports results from stacking changes in commuting zone high-school dropout rates and changes in exposure to Chinese imports within local industries over the periods 1991–99 and 1999–2011. The dependent variable is the annual change in the high-school dropout rate (N = 1,444 = 722 commuting zones × 2 periods). High-school dropout rate is computed from the US Census for 1990 and 2000, and from the ACS for 2011 as a five-year average. In Panel B, each column reports results from regressing changes in commuting zone measures of absolute college mobility on changes in exposure to Chinese imports within local industries over the period 1999–2011. The dependent variable is the annual change in college mobility between cohorts born in 1984 and 1993 (N = 616 commuting zones). College mobility is CZ-level average of college attendance of children with parents at the 25th percentile in the national distribution. The college mobility measure comes from Chetty et al. (2014) and is based on the US tax records. In Panels A and B, the explanatory variable is an employment-weighted average of annualized changes in exposure to Chinese imports within local industries, as detailed in the text. In Columns (3) and (4), the import exposure is instrumented with the alternative instrument constructed from Chinese imports to eight other high-income countries, as in Autor et al. (2013). The commuting zone baseline manufacturing controls are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 period. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. Other baseline controls include population counts, employment-to-population ratios, and region controls for nine regional census divisions. All models include a time trend. Standard errors are clustered by commuting zone.
$*p < 0.10$

$**p < 0.05$
$***p < 0.01$

Table D.3: Log-Log Specification and Baseline Control for Outcome: 2SLS Estimates of Trade Exposure Effects on High-School Dropout Rate 1991–2011.

High-School Dropout Rate	2SLS	
	(1)	(2)
	A. Log-Log Specification	
Commuting zone manufacturing decline	-.865***	-.498***
	(.392)	(.111)
Baseline manufacturing emp. share	–	Yes
Other baseline controls	–	Yes
	B. Baseline Control for High-School Dropout Rate	
Commuting zone import exposure	-.120***	-.397***
	(.029)	(.117)
Baseline Control for High-School Dropout Rate	.039***	.042***
	(.0023)	(.0030)
Baseline manufacturing emp. share	–	Yes
Other baseline controls	–	Yes

Notes: Log-Log Specification and Baseline Control for Outcome. In Panel A, each column reports results from stacking the logarithms of changes in commuting zone high-school dropout rates and declines in manufacturing-to-total employment ratios over the periods 1991–99 and 1999–2011. The dependent variable is the annual change in the high-school dropout rate (N = 1,444 = 722 commuting zones × 2 periods). High-school dropout rate is computed from the US Census for 1990 and 2000, and from the ACS for 2011 as a five-year average. In Panel B, each column reports results from stacking the logarithms of changes in commuting zone high-school dropout rates and declines in manufacturing-to-total employment ratios over the periods 1991–99 and 1999–2011, including controls for the start-of-period high-school dropout rate. In Panels A and B, the manufacturing decline is instrumented with the commuting zone import exposure from China's imports. The instrument is an employment-weighted average of annualized changes in exposure to Chinese imports within local industries, as detailed in the text. The commuting zone baseline controls are computed in 1991 for the 1991–99 period and in 1999 for the 1999–2011 period. Manufacturing employment is computed from the CBP; population data come from the Census Population Estimates. The other baseline controls include population counts, employment-to-population ratios, and region controls for nine regional census divisions. All models include a time trend. Standard errors are clustered by commuting zone.
$*p < 0.10$
$**p < 0.05$
$***p < 0.01$

Table D.4: Correlates of the Intergenerational Effects: Variable Definitions, Part I.

Interaction term	
Segregation and Race	
Fraction Black	Number of individuals who are black alone divided by total population. US Census 2000.
Income Segregation	Rank-Order index estimated at the census-tract level using equation (13) in Reardon (2011); the δ vector is given in Appendix A4 of Reardon's paper. H(pk) is computed for each of the income brackets given in the 2000 census. See Appendix D for further details. US Census 2000.
Segregation of Affluence ($>$p75)	H(p275) estimated following Reardon (2011); we compute H(p) for 16 income groups defined by the 2000 census. We estimate H(p75) using a fourth-order polynomial of the weighted linear regression in equation (12) of Reardon (2011). US Census 2000.
Fraction with Commute $<$ 15 Mins	Number of workers that commute less than 15 minutes to work divided by total number of workers. Sample restricts to workers that are 16 or older and not working at home. US Census 2000.
Income Inequality	
Household Income per Capita	Aggregate household income in the 2000 census divided by the number of people aged 16-64. US Census 2000.
Gini coefficient	Gini coefficient computed using parents of children in the core sample, with income topcoded at $100 million in 2012 dollars. Tax Records, Core Sample.
Fraction Middle Class	Fraction of parents (in the core sample) whose income falls between the 25th and 75th percentile of the national parent income distribution. Tax Records, Core Sample.
K-12 Education	
School Expenditure per Student.	Average expenditures per student in public schools. NCES CCD 1996-1997 Financial Survey.
Student Teacher Ratio	Average student-teacher ratio in public schools. NCES CCD 1996-1997 Universe Survey
Test Score Percentile	Residual from a regression of mean math and English standardized test scores on household income per capita in 2000. George Bush Global Report Card.

Notes: These covariates are compiled by Chetty et al. (2014). The descriptions come from that source. See the reference for further details.

Table D.5: Correlates of the Intergenerational Effects: Variable Definitions, Part II.

Interaction term	
College	
Number of Colleges per Cap.	Number of Title IV, degree offering institutions per capita. IPEDS 2000
College Tuition	Mean in-state tuition and fees for first-time, full-time undergraduates. IPEDS 2000.
College Graduation Rate	Residual from a regression of graduation rate (the share of undergraduate students that complete their degree in 150% of normal time) on household income per capita in 2000. IPEDS 2009.
Social Capital	
Social Capital Index	Standardized index combining measures of voter turnout rates, the fraction of people who return their census forms, and measures of participation in community organizations. Rupasingha and Goetz (2008).
Fraction Religious	Share of religious adherents. Association of Religion Data Archives
Violent Crime Rate	Number of arrests for serious violent crimes per capita. Uniform Crime Reports.
Local Labor Market	
Teenage (14-16) LFP	Fraction of children in birth cohorts 1985-1987 who received a W2 (i.e. had positive wage earnings) in any of the tax years when they were age 14-16. Tax Records, Extended Sample.

Notes: These covariates are compiled by Chetty et al. (2014). The descriptions come from that source. See the reference for further details.

Bibliography

Acemoglu, Daron, "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *The Quarterly Journal of Economics*, 1998, *113* (4), 1055–1089.

_ , "Directed Technical Change," *The Review of Economic Studies*, 2002, *69* (4), 781–809.

_ , "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature*, 2002, *40* (1), 7–72.

_ **and David H Autor**, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in David Card and Orley Ashenfelter, eds., *Handbook of Labor Economics*, Vol. 4B 2011, pp. 1043–1171.

_ **and Pascual Restrepo**, "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment," *American Economic Review*, 2018, *108* (6), 1488–1542.

_ **and** _ , "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

_ **and** _ , "Demographics and Automation," *The Review of Economic Studies*, 2021.

_ **, Andrea Manera, and Pascual Restrepo**, "Does the U.S. Tax Code Favor Automation?," *Brookings Papers on Economic Activity*, 2020, (Spring), 231–300.

_ **, Claire Lelarge, and Pascual Restrepo**, "Competing with Robots: Firm-Level Evidence from France," *AEA Papers and Proceedings*, 2020, *110*, 383–88.

_ **, David H Autor, David Dorn, Gordon H Hanson, and Brendan Price**, "Import Competition and the Great US Employment Sag of the 2000s," *Journal of Labor Economics*, 2016, *34* (S1), S141–S198.

_ **, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr**, "Innovation, Reallocation and Growth," *American Economic Review*, 2018, *108* (11), 3450–91.

Aghion, Philippe and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica*, 1992, *60* (2), 323–351.

_ **, Celine Antonin, Simon Bunel, and Xavier Jaravel**, "What are the Labor and Product Market Effects of Automation? New Evidence from France," *Working Paper*, 2020.

_ **, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt**, "Competition and Innovation: an Inverted-U Relationship," *The Quarterly Journal of Economics*, 2005, *120* (2), 701–728.

Akcigit, Ufuk and William R Kerr, "Growth through Heterogeneous Innovations," *Journal of Political Economy*, 2018, *126* (4), 1374–1443.

Akee, Randall K.Q., William Copeland, and Ej Costello, "Parents' Incomes and Children's Outcomes: A Quasi-Experiment," *American Economic Journal: Applied Economics*, 2010, *2* (1), 86–115.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad, "The Skill Complementarity of Broadband Internet," *The Quarterly Journal of Economics*, 2015, *130* (4), 1781–1824.

Alexopoulos, Michelle, "Read All about it!! What Happens Following a Technology Shock?," *American Economic Review*, 2011, *101* (4), 1144–79.

Almlund, Mathilde, Angela Lee Duckworth, James Heckman, and Tim Kautz, "Personality Psychology and Economics," in "Handbook of the Economics of Education," Vol. 4, Elsevier, 2011, pp. 1–181.

Altonji, Joseph G, Peter Arcidiacono, and Arnaud Maurel, "The analysis of field choice in college and graduate school: Determinants and wage effects," in "Handbook of the Economics of Education," Vol. 5, Elsevier, 2016, pp. 305–396.

Ananat, Elizabeth O, Anna Gassman, Dania V Francis, and M Christina, "Linking job loss, inequality, mental health, and education," *Science*, 2017, *356* (6343), 1127–1128.

_ , Anna Gassman-Pines, and Christina M Gibson-Davis., "The Effects of Local Employment Losses on Children's Educational Achievement," in Greg J Duncan and Richard J. Murnane, eds., *Whither Opportunity?: Rising Inequality, Schools, and Children's Life Chances*, Russell Sage Foundation, 2011, pp. 299–314.

Angrist, Joshua D., "Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants," *Econometrica*, 1998, *66* (2), 249–288.

Angrist, Joshua D and Jorn-Steffen Pischke, *Mostly Harmless Econometrics*, Princeton University Press, 2009.

Atalay, Enghin, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum, "The Evolution of Work in the United States," *American Economic Journal: Applied Economics*, 2020, *12* (2), 1–34.

Atkin, David, "Endogenous Skill Acquisition and Export Manufacturing in Mexico," *American Economic Review*, 2016, *106* (8), 2046–85.

Autor, David H, "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 2015, *29* (3), 3–30.

_ and David Dorn, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 2013, *103* (5), 1553–1597.

_ , Anna Salomons, and Bryan Seegmiller, "New Frontiers: The Origins and Content of New Work, 1940-2018," *Working Paper*, 2021.

_ , David Dorn, and Gordon H Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

_ , _ , and _ , "Untangling Trade and Technology: Evidence from Local Labour Markets," *Economic Journal*, 2015, *125* (584), 621–646.

_ , _ , and _ , "The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade," *Annual Review of Economics*, 2016, *8* (1), 205–240.

_ , _ , and _ , "When Work Disappears: Manufacturing Decline and the Falling Marriage-Market Value of Men," *NBER Working Paper No. 23173*, 2017.

_ , _ , _ , and Jae Song, "Trade Adjustment: Worker-Level Evidence," *The Quarterly Journal of Economics*, 2014, *129* (4), 1799–1860.

_ , Frank Levy, and Richard J Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

_ , Lawrence F Katz, and Alan B Krueger, "Computing Inequality: Have Computers Changed the Labor Market?," *The Quarterly Journal of Economics*, 1998, *113* (4), 1169–1213.

Balsvik, Ragnhild, Sissel Jensen, and Kjell G. Salvanes, "Made in China, sold in Norway: Local labor market effects of an import shock," *Journal of Public Economics*, 2013, *127*, 137–144.

Bartel, Ann, Casey Ichniowski, and Kathryn Shaw, "How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills," *The Quarterly Journal of Economics*, 2007, *122* (4), 1721–1758.

Beaudry, Paul, Mark Doms, and Ethan Lewis, "Should the Personal Computer Be Considered a Technological Revolution? Evidence from U.S. Metropolitan Areas," *Journal of Political Economy*, 2010, *118* (5), 988–1036.

Becker, Gary, *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, University of Chicago Press, 1964.

Becker, Sascha O., Peter H. Egger, and Maximilian Von Ehrlich, "Going NUTS: The effect of EU Structural Funds on regional performance," *Journal of Public Economics*, 2010, *94* (9-10), 578–590.

Berger, Suzanne, *Making in America: From Innovation to Market*, Cambridge, MA: MIT Press, 2013.

__ , "Manufacturing in America: A View from the Field," *MIT Task Force on the Work of the Future Research Brief 16*, 2020.

__ **and MIT Industrial Performance Center**, *How We Compete: What Companies Around the World Are Doing to Make it in Today's Global Economy*, New York: Currency Doubleday, 2005.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F Katz, and Christopher Palmer, "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice," *NBER Working Paper 26164*, 2019.

Berman, E., J. Bound, and Z. Griliches, "Changes in the Demand for Skilled Labor within U. S. Manufacturing: Evidence from the Annual Survey of Manufactures," *The Quarterly Journal of Economics*, 1994, *109* (2), 367–397.

Bernard, Andrew B. and J. Bradford Jensen, "Exporters, skill upgrading, and the wage gap," *Journal of International Economics*, 1997, *42* (1), 3–31.

__ **, Stephen J. Redding, and Peter K. Schott**, "Multiple-Product Firms and Product Switching," *American Economic Review*, 2010, *100* (1), 70–97.

__ **,** __ **, and** __ , "Multiproduct Firms and Trade Liberalization," *The Quarterly Journal of Economics*, 2011, *126* (3), 1271–1318.

Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge, "What Happens to Workers at Firms that Automate?," *Working Paper*, 2020.

Bils, Mark and Peter J. Klenow, "The Acceleration of Variety Growth," *American Economic Review*, 2001, *91* (2), 274–280.

Bird, Steven, Ewan Klein, and Edward Loper, *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*, " O'Reilly Media, Inc.", 2009.

Black, Dan A, Terra G Mckinnish, and Seth G Sanders, "Tight Labor Markets and the Demand for Education: Evidence from the Coal Boom and Bust," *Industrial and Labor Relations Review*, 2005, *59* (1), 3–16.

Blanchard, Olivier J and Lawrence F Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, pp. 1–75.

Bloom, Nick, Mirko Draca, and John Van Reenen, "Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity," *Review of Economic Studies*, 2016, *83* (1), 87–117.

Bojanowski, Piotr, Edouard Grave, Armand Joulin, and Tomas Mikolov, "Enriching Word Vectors with Subword Information," *arXiv preprint arXiv:1607.04606*, 2016.

Boler, Esther, "Technology-Skill Complementarity in a Globalized World," *Working Paper*, 2015.

Bonfiglioli, Alessandra, Rosario Crino, Harald Fadinger, and Gino Gancia, "Robot Imports and Firm-Level Outcomes," *Working Paper*, 2020.

Boone, Jan, "Technological Progress, Downsizing and Unemployment," *The Economic Journal*, 2000, *110* (465), 581–600.

Borghans, Lex, Bart HH Golsteyn, James J Heckman, and John Eric Humphries, "What grades and achievement tests measure," *Proceedings of the National Academy of Sciences*, 2016, *113* (47), 13354–13359.

Boyd, Melofy and Stefanie DeLuca, “Fieldwork with in-depth interviews: How to get strangers in the city to tell you their stories,” in “Methods in Social Epidemiology,” John Wiley & Sons, 2017.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M Hitt, “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence,” *The Quarterly Journal of Economics*, 2002, *117* (1), 339–376.

Brynjolfsson, Erik and Andrew McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, WW Norton & Company, 2014.

Buri, Riku, “Hyppyja tukitasossa: Investointitukien vaikutusten ekonometrinen selvitys,” *Aalto University MA Thesis*, 2017.

Bursztyn, Leonardo, Georgy Egorov, and Robert Jensen, “Cool to be Smart or Smart to be Cool? Understanding Peer Pressure in Education,” *The Review of Economic Studies*, 2019, *86* (4), 1487–1526.

Bustos, Paula, “Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms,” *American Economic Review*, 2011, *101* (1), 304–340.

Caballero, Ricardo J. and Eduardo M. R. A. Engel, “Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach,” *Econometrica*, 1999, *67* (4), 783–826.

Card, David and John E DiNardo, “Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles,” *Journal of Labor Economics*, 2002, *20* (4), 733–783.

Caroli, Eve and John Van Reenen, “Skill-Biased Organizational Change? Evidence from A Panel of British and French Establishments,” *The Quarterly Journal of Economics*, 2001, *116* (4), 1449–1492.

Cattaneo, Matias D., Michael Jansson, and Xinwei Ma, “Manipulation testing based on density discontinuity,” *The Stata Journal*, 2018, *18* (1), 234–261.

Caves, Douglas W., Laurits R. Christensen, and Joseph A. Swanson, “Productivity Growth, Scale Economies, and Capacity Utilization in U.S. Railroads, 1955-74,” *The American Economic Review*, 1981, *71* (5), 994–1002.

Cerqua, Augusto and Guido Pellegrini, “Do subsidies to private capital boost firms’ growth? A multiple regression discontinuity design approach,” *Journal of Public Economics*, 2014, *109*, 114–126.

Chetty, Raj and Nathaniel Hendren, “The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects,” *NBER Working Paper 23001*, 2017.

_ and _ , “The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates,” *NBER Working Paper 230002*, 2017.

_ , _ , Patrick Kline, and Emmanuel Saez, “Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States,” *Quarterly Journal of Economics*, 2014, *129* (4), 1553–1623.

Ciccone, Antonio and Elias Papaioannou, “Human Capital, the Structure of Production, and Growth,” *The Review of Economics and Statistics*, 2009, *91* (1), 66–82.

Cooper, Russell, John Haltiwanger, and Laura Power, “Machine Replacement and the Business Cycle: Lumps and Bumps,” *American Economic Review*, 1999, *89* (4), 921–946.

Corr, Philip J and Gerald Matthews, *The Cambridge Handbook of Personality Psychology*, Cambridge University Press, 2020.

Costinot, Arnaud and Ivan Werning, “Robots, Trade, and Luddism: A Sufficient Statistics Approach to Optimal Technology Regulation,” *Working Paper*, 2020.

Criscuolo, Chiara, Ralf Martin, Henry G. Overman, and John Van Reenen, “Some Causal Effects of an Industrial Policy,” *American Economic Review*, 2019, *109* (1), 48–85.

Cunha, Flavio and James Heckman, "The Technology of Skill Formation," *American Economic Review*, 2007, *97* (2), 31–47.

Curtis, E. Mark, Daniel G. Garrett, Eric C. Ohrn, Kevin A. Roberts, and Juan Carlos Suarez Serrato, "Capital Investment and Labor Demand," *NBER Working Paper 29485*, 2021.

Dahlberg, Matz, Linna Marten, and Bjorn Ockert, "Who recovers from a job loss? The importance of cognitive and non-cognitive skills," *Working Paper*, 2021.

Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner, "The Adjustment of Labor Markets to Robots," *Journal of the European Economic Association*, 2021.

Davis, Steven J. and Till von Wachter, "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, 2011, *2011* (2), 1–72.

Dechezlepretre, Antoine, David Hemous, Morten Olsen, and Carlo Zanella, "Induced Automation: Evidence from Firm-level Patent Data," *Working Paper*, 2021.

Deiana, Claudio, "The Bitter Side of Trade Shocks: Local Labour Market Conditions and Crime in the US," *Working Paper*, 2016.

Deming, David J, "The growing importance of social skills in the labor market," *The Quarterly Journal of Economics*, 2017, *132* (4), 1593–1640.

Dertouzos, Michael L., Robert M. Solow, and Richard K. Lester, *Made in America: Regaining the Productive Edge*, Cambridge, MA: MIT Press, 1989.

Dix-Carneiro, Rafael and Brian K Kovak, "Trade Liberalization and Regional Dynamics," *American Economic Review (forthcoming)*, 2017.

Dixit, Avinash K. and Joseph E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity," *The American Economic Review*, 1977, *67* (3), 297–308.

Dixon, Jay, Bryan Hong, and Lynn Wu, "The Robot Revolution: Managerial and Employment Consequences for Firms," *Management Science*, 2021.

Doms, Mark and Timothy Dunne, "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics*, 1998, *1* (2), 409–429.

_ , _ , and Kenneth R Troske, "Workers, Wages, and Technology," *The Quarterly Journal of Economics*, 1997, *112* (1), 253–290.

Dorn, David, "Essays on Inequality, Spatial Interaction, and the Demand for Skills," *Dissertation University of St. Gallen No. 3613.*, 2009.

Economist, "Left behind: How to help places hurt by globalization," 2017, *October.*

Edin, Per-Anders, Peter Fredriksson, Martin Nybom, and Bjorn Ockert, "The Rising Return to Non-cognitive Skill," *American Economic Journal: Applied Economics*, 2021.

Edmonds, Eric V, Nina Pavcnik, and Petia Topalova, "Trade Adjustment and Human Capital Investments: Evidence from Indian Tariff Reform," *American Economic Journal: Applied Economics*, 2010, *2* (4), 42–75.

Eggleston, Karen, Yong Suk Lee, and Toshiaki Iizuka, "Robots and Labor in the Service Sector: Evidence from Nursing Homes," *NBER Working Paper 28322*, 2021.

Einio, Elias, "R&D subsidies and company performance: Evidence from geographic variation in government funding based on the ERDF population-density rule," *The Review of Economics and Statistics*, 2014, *96* (4), 710–728.

Ellison, Glenn, Edward L Glaeser, and William R Kerr, "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns," *American Economic Review*, 2010, *100* (June), 1195–1213.

Ericson, Richard and Ariel Pakes, "Markov-Perfect Industry Dynamics: A Framework for Empirical work," *The Review of Economic Studies*, 1995, *62* (1), 53–82.

Evans, W and W Kim, "The impact of local labor market conditions on the demand for education: evidence from Indian casinos," *Working Paper*, 2008.

Feigenbaum, James and Daniel P. Gross, "Automation and the Future of Young Workers: Evidence from Telephone Operation in the Early 20th Century," *NBER Working Paper 28061*, 2021.

Feler, Leo and Mine Z. Senses, "Trade shocks and the provision of local public goods," *American Economic Journal: Economic Policy*, 2017, *9* (4), 101–143.

Fieler, Ana Cecilia and Ann Harrison, "Escaping Import Competition in China," *NBER Working Paper 24527*, 2018.

Flam, Harry and Elhanan Helpman, "Industrial policy under monopolistic competition," *Journal of International Economics*, 1987, *22* (1), 79–102.

Florida, Richard, *The New Urban Crisis: How Our Cities Are Increasing Inequality, Deepening Segregation, and Failing the Middle Class - and What We Can Do About It*, Basic Books, 2017.

Ford, Henry, *My Life and Work*, Garden City, N.Y.: Doubleday, Page & Co., 1922.

Fredriksson, Peter, Lena Hensvik, and Oskar Nordström Skans, "Mismatch of talent: Evidence on match quality, entry wages, and job mobility," *American Economic Review*, 2018, *108* (11), 3303–38.

Gaggl, Paul and Greg C. Wright, "A short-run view of what computers do: Evidence from a UK tax incentive," *American Economic Journal: Applied Economics*, 2017, *9* (3), 262–294.

Ganong, Peter and Daniel Shoag, "Why Has Regional Income Convergence in the U.S. Declined?," *Journal of Urban Economics*, 2017, *102* (November), 76–90.

Gardner, Howard, *Frames of Mind: The Theory of Multiple Intelligences*, Basic Books, 1983.

Gathmann, Christina and Uta Schönberg, "How general is human capital? A task-based approach," *Journal of Labor Economics*, 2010, *28* (1), 1–49.

Gelman, Andrew and Guido Imbens, "Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs," *Journal of Business & Economic Statistics*, 2019, *37* (3), 447–456.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, "Text as Data," *Journal of Economic Literature*, 2019, *57* (3), 535–574.

Genz, Sabrina, Terry Gregory, Markus Janser, Florian Lehmer, and Britta Matthes, "How Do Workers Adjust When Firms Adopt New Technologies?," *Working Paper*, 2021.

Giorcelli, Michela, "The Long-Term Effects of Management and Technology Transfers," *American Economic Review*, 2019, *109* (1), 121–152.

Glaeser, Edward and Joseph Gyourko, "Urban Decline and Durable Housing," *Journal of Political Economy*, 2005, *113* (2), 345–000.

Glaeser, Edward L. and Joshua D. Gottlieb, "The Economics of Place-Making Policies," *Brookings Papers on Economic Activity*, 2008, (Spring).

Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova, "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India," *The Quarterly Journal of Economics*, 2010, *125* (4), 1727–1767.

Goldin, Claudia and Lawrence F Katz, "Education and Income in the Early Twentieth Century: Evidence from the Prairies," *Journal of Economic History*, 2000, *60* (03), 782–818.

— **and Lawrence Katz**, *The Race Between Education and Technology*, Belknap Press for Harvard University Press, 2008.

Gollop, Frank M. and James L. Monahan, "A Generalized Index of Diversification: Trends in U.S. Manufacturing," *The Review of Economics and Statistics*, 1991, *73* (2), 318–330.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021.

Graetz, Georg and Guy Michaels, "Robots at Work," *The Review of Economics and Statistics*, 2018, *100* (5), 753–768.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti, "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings," *Journal of Political Economy*, 2010, *118* (3), 536–598.

Griliches, Zvi, "Capital-Skill Complementarity," *The Review of Economics and Statistics*, 1969, *51* (4), 465–468.

Groes, Fane, Philipp Kircher, and Iourii Manovskii, "The U-shapes of occupational mobility," *The Review of Economic Studies*, 2015, *82* (2), 659–692.

Grossman, Gene M. and Elhanan Helpman, "Quality Ladders in the Theory of Growth," *The Review of Economic Studies*, 1991, *58* (1), 43–61.

Gruber, Jonathan and Simon Johnson, *Jump-Starting America: How Breakthrough Science Can Revive Economic Growth and the American Dream*, New York: Public Affairs, 2019.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles, "Should Robots Be Taxed?," *The Review of Economic Studies*, 2021.

Guvenen, Fatih, Burhan Kuruscu, Satoshi Tanaka, and David Wiczer, "Multidimensional skill mismatch," *American Economic Journal: Macroeconomics*, 2020, *12* (1), 210–44.

Hamermesh, Daniel S., "Labor Demand and the Structure of Adjustment Costs," *The American Economic Review*, 1989, *79* (4), 674–689.

Hanson, Gordon H, "The Rise of Middle Kingdoms: Emerging Economies in Global Trade," *Journal of Economic Perspectives*, 2012, *26* (2), 41–64.

Harrison, Rupert, Jordi Jaumandreu, Jacques Mairesse, and Bettina Peters, "Does innovation stimulate employment? A firm-level analysis using comparable micro-data from four European countries," *International Journal of Industrial Organization*, 2014, *35* (July), 29–43.

Hausmann, Ricardo, Jason Hwang, and Dani Rodrik, "What you export matters," *Journal of Economic Growth*, 2007, *12* (1), 1–25.

Hawkins, William, Ryan Michaels, and Jiyoon Oh, "The Joint Dynamics of Capital and Employment at the Plant Level," *Working Paper*, 2015.

Heckman, James J, Jora Stixrud, and Sergio Urzua, "The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior," *Journal of Labor Economics*, 2006, *24* (3), 411–482.

Hemous, David and Morten Olsen, "The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality," *American Economic Journal: Macroeconomics*, 2021.

Hirano, Keisuke, Guido W. Imbens, and Geert Ridder, "Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score," *Econometrica*, 2003, *71* (4), 1161–1189.

Hopenhayn, Hugo A., "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 1992, pp. 1127–1150.

Howell, Sabrina T., "Financing Innovation: Evidence from R&D Grants," *American Economic Review*, 2017, *107* (4), 1136–64.

_ , Jason Rathje, John Van Reenen, and Jun Wong, "Opening up Military Innovation: Causal Effects of 'Bottom-Up' Reforms to US Defense Research," *NBER Working Paper 28700*, 2021.

Hoynes, Hilary, Douglas L Miller, and Jessamyn Schaller, "Who Suffers During Recessions?," *Journal of Economic perspectives*, 2012, *26* (3), 27–48.

Humlum, Anders, "Robot Adoption and Labor Market Dynamics," *Working Paper*, 2019.

Huttunen, Kristiina, Jarle Moen, and Kjell G. Salvanes, "How Destructive Is Creative Destruction? Effects of Job Loss on Job Mobility, Withdrawal and Income," *Journal of the European Economic Association*, 2011, *9* (5), 840–870.

Iacus, Stefano M., Gary King, and Giuseppe Porro, "Causal Inference without Balance Checking: Coarsened Exact Matching," *Political Analysis*, 2012, *20* (1), 1–24.

Imbens, Guido W and Jeffrey M Wooldridge, "Recent Developments in the Econometrics of Program Evaluation," *Journal of Economic Literature*, 2009, *47* (1), 5–86.

Izadi, Ramin and Joonas Tuhkuri, "Psychological Traits and Adaptation in the Labor Market," *Working Paper*, 2021.

_ and _ , "School vs. Action-Oriented Personalities in the Labor Market," *Working Paper*, 2021.

Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan, "Earnings Losses of Displaced Workers," *The American Economic Review*, 1993, *83* (4), 685–709.

Jensen, Robert, "Do labor market opportunities affect young women's work and family decisions? Experimental evidence from India," *The Quarterly Journal of Economics*, 2012, *127* (2), 753–792.

Jokela, Markus, Tuomas Pekkarinen, Matti Sarvimäki, Marko Terviö, and Roope Uusitalo, "Secular rise in economically valuable personality traits," *Proceedings of the National Academy of Sciences*, 2017, *114* (25), 6527–6532.

Jussim, Lee, Jarret T Crawford, Stephanie M Anglin, John R Chambers, Sean T Stevens, and Florette Cohen, "Stereotype accuracy: One of the largest and most replicable effects in all of social psychology," *Handbook of prejudice, stereotyping, and discrimination*, 2016, *2*, 31–63.

Kaila, Martti, Emily Nix, and Krista Riukula, "Disparate Impacts of Job Loss by Parental Income and Implications for Intergenerational Mobility," *Working Paper*, 2021.

Katz, Lawrence F and Kevin Murphy, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 1992, *107* (1), 35–78.

Kauhanen, Antti and Krista Riukula, "The Costs of Job Loss and Task Usage," *ETLA Working Papers No 73*, 2019.

Kekkonen, Urho, *Onko maallamme malttia vaurastua?*, Otava, 1952.

Kenney, Martin and Richard Florida, *Beyond Mass Production*, Oxford University Press, 1993.

Kerr, Sari, Terhi Maczulskij, and Mika Maliranta, "Within and between firm trends in job polarization: the roles of globalization and technology," *Journal of Economic Geography*, 2020, *20* (4), 1003–1039.

Keynes, John Maynard, "Economic Possibilities for Our Grandchildren," in "Essays in Persuasion," London: Macmillan, 1931.

Klette, Tor Jakob and Samuel Kortum, "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 2004, *112* (5), 986–1018.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar, "Who Profits from Patents? Rent-Sharing at Innovative Firms," *The Quarterly Journal of Economics*, 2019, *134* (3), 1343–1404.

Koch, Michael, Ilya Manuylov, and Marcel Smolka, "Robots and Firms," *The Economic Journal*, 2021, *131* (638), 2553–2584.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller, "Technological Change and Occupations over the Long Run," *Working Paper*, 2020.

Koren, Miklos, Marton Csillag, and Janos Kollo, "Machines and Machinists: Incremental Technical Change and Wage Inequality," *Working Paper*, 2020.

Kremer, Michael, “The O-ring Theory of Economic Development,” *The Quarterly Journal of Economics*, 1993, *108* (3), 551–575.

Krolikowski, Pawel, “Choosing a control group for displaced workers,” *ILR Review*, 2018, *71* (5), 1232–1254.

Krusell, Per, Lee E. Ohanian, Jose-Victor Rios-Rull, and Giovanni L. Violante, “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 2000, *68* (5), 1029–1053.

Kugler, M. and E. Verhoogen, “Prices, Plant Size, and Product Quality,” *The Review of Economic Studies*, 2012, *79* (1), 307–339.

Lachowska, Marta, Alexandre Mas, and Stephen A Woodbury, “Sources of Displaced Workers' Long-Term Earnings Losses,” *American Economic Review*, 2020, *110* (10), 3231–66.

Lane, Nathan, “The New Empirics of Industrial Policy,” *Journal of Industry, Competition, and Trade*, 2020, *20* (2), 209–234.

_ , “Manufacturing Revolutions: Industrial Policy and Industrialization in South Korea,” *Working Paper*, 2021.

Lavecchia, A.M., H. Liu, and P. Oreopoulos, “Chapter 1 - Behavioral Economics of Education: Progress and Possibilities,” in Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, eds., *Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, eds.*, Vol. 5 of *Handbook of the Economics of Education*, Elsevier, 2016, pp. 1–74.

Leiponen, Aija, “Skills and innovation,” *International Journal of Industrial Organization*, 2005, *23* (5-6), 303–323.

Levine, Ross and Yona Rubinstein, “Smart and Illicit: Who Becomes an Entrepreneur and Do They Earn More?,” *The Quarterly Journal of Economics*, 2017, *132* (2), 963–1018.

Levinsohn, James and Amil Petrin, “Estimating Production Functions Using Inputs to Control for Unobservables,” *The Review of Economic Studies*, 2003, *70* (2), 317–341.

Lewis, Ethan, “Immigration, Skill Mix, and Capital Skill Complementarity,” *The Quarterly Journal of Economics*, 2011, *126* (2), 1029–1069.

Lileeva, Alla and Daniel Trefler, “Improved access to foreign markets raises plant-level productivity...for some plants,” *The Quarterly Journal of Economics*, 2010, *125* (3), 1051–1099.

Lindenlaub, Ilse, “Sorting multidimensional types: Theory and application,” *The Review of Economic Studies*, 2017, *84* (2), 718–789.

Lindner, Attila, Balazs Murakozy, Balazs Reizer, and Ragnhild Schreiner, “Firm-level Technological Change and Skill Demand,” *Working Paper*, 2021.

Lindqvist, Erik and Roine Vestman, “The labor market returns to cognitive and noncognitive ability: Evidence from the Swedish enlistment,” *American Economic Journal: Applied Economics*, 2011, *3* (1), 101–28.

Lise, Jeremy and Fabien Postel-Vinay, “Multidimensional skills, sorting, and human capital accumulation,” *American Economic Review*, 2020, *110* (8), 2328–76.

Lleras-Muney, Adriana, Matthew Miller, Shuyang Sheng, and Veronica T Sovero, “Party On: The Labor Market Returns to Social Networks and Socializing,” *NBER Working Paper No. 27337*, 2020.

Manera, Andrea and Martina Uccioli, “Employment Protection and the Direction of Technology Adoption,” *Working Paper*, 2021.

Mann, Katja and Lukas Puttmann, “Benign Effects of Automation: New Evidence from Patent Texts,” *The Review of Economics and Statistics*, 2021.

Marx, Karl, *Capital. Volume I: The Process of Production of Capital* 1867.

Matsuyama, Kiminori, “Beyond Icebergs: Towards a Theory of Biased Globalization,” *The Review of Economic Studies*, 2007, *74* (1), 237–253.

McCrary, Justin, "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of Econometrics*, 2008, *142* (2), 698–714.

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

_ **and Stephen J. Redding**, "Heterogeneous Firms and Trade," in "Handbook of International Economics," Vol. 4, Elsevier, 2014, pp. 1–54.

Mian, Atif and Amir Sufi, "What Explains the 2007-2009 Drop in Employment?," *Econometrica*, 2014, *82* (6), 2197–2223.

Michaels, Guy, Ashwini Natraj, and John Van Reenen, "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years," *The Review of Economics and Statistics*, 2014, *96* (1), 60–77.

Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean, "Efficient Estimation of Word Representations in Vector Space," *arXiv preprint arXiv:1301.3781*, 2013.

Milgrom, Paul and John Roberts, "The Economics of Modern Manufacturing: Technology, Strategy, and Organization," *The American Economic Review*, 1990, pp. 511–528.

Mitrunen, Matti, "Structural Change and Intergenerational Mobility: Evidence from the Finnish War Reparations," *Working Paper*, 2021.

Moretti, Enrico, *The New Geography of Jobs*, Houghton Mifflin Harcourt, 2012.

Mõttus, René, "Towards more rigorous personality trait–outcome research," *European Journal of Personality*, 2016, *30* (4), 292–303.

Mozer, Reagan, Luke Miratrix, Aaron Russell Kaufman, and L. Jason Anastasopoulos, "Matching with Text Data: An Experimental Evaluation of Methods for Matching Documents and of Measuring Match Quality," *Political Analysis*, 2020, *28* (4), 445–468.

Munshi, Kaivan and Mark Rosenzweig, "Traditional institutions meet the modern world: Caste, gender, and schooling choice in a globalizing economy," *American Economic Review*, 2006, *96* (4), 1225–1252.

Naughton, Barry J., *The Chinese Economy*, MIT Press, 2006.

Nelson, Richard R. and Edmund S Phelps, "Investment in Humans, Technological Diffusion, and Economic Growth," *The American Economic Review*, 1966, *56* (1/2), 69–75.

Nilsen, Oivind A. and Fabio Schiantarelli, "Zeros and Lumps in Investment: Empirical Evidence on Irreversibilities and Nonconvexities," *The Review of Economics and Statistics*, 2003, *85* (4), 1021–1037.

_ **, Arvid Raknerud, Marina Rybalka, and Terje Skjerpen**, "Lumpy investments, factor adjustments, and labour productivity," *Oxford Economic Papers*, 2009, *61* (1), 104–127.

Notowidigdo, Matthew J, "The Incidence of Local Labor Demand Shocks," *NBER Working Paper No. 17167*, 2011.

Nyman, Kai, "Varusmiesten johtajavalintojen luotettavuus," *Publication Series 1, National Defense University, Department of Behavioral Sciences, Helsinki, Finland*, 2007.

Oberfield, Ezra and Devesh Raval, "Micro Data and Macro Technology," *Econometrica*, 2021, *89* (2), 703–732.

Ohlsbom, Roope and Mika Maliranta, "Management Practices and Allocation of Employment: Evidence from Finnish Manufacturing," *International Journal of the Economics of Business*, 2021, *28* (1), 115–138.

Olley, Steven and Ariel Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," 1992.

Oreopoulos, Philip, Marianne Page, and Ann Huff Stevens, "The Intergenerational Effects of Worker Displacement," *Journal of Labor Economics*, 2008, *26* (3), 455–483.

Oster, Emily and Bryce Millett Steinberg, "Do IT service centers promote school enrollment? Evidence from India," *Journal of Development Economics*, 2013, *104*, 123–135.

Papageorge, Nicholas W, Victor Ronda, and Yu Zheng, "The Economic Value of Breaking Bad: Misbehavior, Schooling and the Labor Market," *NBER Working Paper 25602*, 2019.

Paunonen, Sampo V and Michael C Ashton, "Big five factors and facets and the prediction of behavior.," *Journal of personality and social psychology*, 2001, *81* (3), 524.

Pavcnik, Nina, "What explains skill upgrading in less developed countries?," *Journal of Development Economics*, 2003, *71* (2), 311–328.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, 2011, *12*, 2825–2830.

Pellegrini, Guido and Teo Muccigrosso, "Do subsidized new firms survive longer? Evidence from a counterfactual approach," *Regional Studies*, 2017, *51* (10), 1483–1493.

Pennington, Jeffrey, Richard Socher, and Christopher D Manning, "Glove: Global Vectors for Word Representation," in "Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)" 2014, pp. 1532–1543.

Pierce, Justin R. and Peter K. Schott, "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, 2016, *106* (7), 1632–1662.

_ and _ , "Trade Liberalization and Mortality: Evidence from U.S. Counties," *Working Paper*, 2016.

Pierce, Justin R and Peter Schott, "A concordance between ten-digit U.S. Harmonized System codes and SIC/NAICS product classes and industries," *Journal of Economic and Social Measurement*, 2012, *37* (1-2), 61–96.

Piore, Michael J, "Qualitative Research Techniques in Economics," *Administrative Science Quarterly*, 1979, *24* (4), 560–569.

_ , "Corporate Reform in American Manufacturing and the Challenge to Economic Theory," in "Information Technology and the Corporation of the 1990s: Research Studies," Oxford University Press, 1994.

_ , "Qualitative research: Does it fit in economics?," *European Management Review*, 2006, *3* (1), 17–23.

_ and Charles Sabel, *The Second Industrial Divide: Possibilities for Prosperity*, New York, NY: Basic Books, 1984.

Porter, Michael E., *The Competitive Advantage: Creating and Sustaining Superior Performance*, NY: Free Press, 1985.

Putnam, Robert D, "Bowling alone: Americas's declining social capital," *Journal of Democracy*, 1995, *6*, 65–78.

Ratcliffe, Michael, Charlynn Burd, Kelly Holder, and Alison Fields, "Defining Rural at the U.S. Census Bureau," *US Census*, 2016, (December).

Rauch, James E., "Networks versus markets in international trade," *Journal of International Economics*, 1999, *48* (1), 7–35.

Raven, John C and J H Court, *Raven's progressive matrices*, Los Angeles, CA: Western Psychological Services, 1938.

Rege, Mari, Kjetil Telle, and Mark Votruba, "Parental job loss and children's school performance," *Review of Economic Studies*, 2011, *78* (4), 1462–1489.

Restrepo, Pascual and Joachim Hubmer, "Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital-Labor Substitution," *Working Paper*, 2021.

Roberts, Margaret E., Brandon M. Stewart, and Richard A. Nielsen, "Adjusting for confounding with text matching," *American Journal of Political Science*, 2020, *64* (4), 887–903.

Rodrik, Dani, "Industrial Policy for the Twenty-First Century," in "One Economics, Many Recipes: Globalization, Institutions, and Economic Growth," Princeton University Press, 2007.

Romer, Christina D. and David H. Romer, "A new measure of monetary shocks: Derivation and implications," *American Economic Review*, 2004, *94* (4), 1055–1084.

Romer, Paul M., "Growth Based on Increasing Returns Due to Specialization," *The American Economic Review*, 1987, *77* (2), 56–62.

— , "Endogenous Technological Change," *Journal of political Economy*, 1990, *98* (5, Part 2), S71–S102.

Rosenbaum, Paul R. and Donald B. Rubin, "The central role of the propensity score in observational studies for causal effects," *Biometrika*, 1983, *70* (1), 41–55.

Ross, Lee and Richard E Nisbett, *The Person and the Situation: Perspectives of Social Psychology*, Mcgraw-Hill Book Company, 1991.

Roy, Andrew Donald, "Some thoughts on the distribution of earnings," *Oxford Economic Papers*, 1951, *3* (2), 135–146.

Saint-Paul, Gilles, "Employment protection, international specialization, and innovation," *European Economic Review*, 2002, *46* (2), 375–395.

Salton, Gerard and Christopher Buckley, "Term-Weighting Approaches in Automatic Text Retrieval," *Information Processing & Management*, 1988, *24* (5), 513–523.

Sanders, Carl, "Reading skills and earnings: Why do doing words good hurt you're wages," *Working Paper*, 2015.

Schmieder, J, Till von Wachter, and Jörg Heining, "The Costs of Job Displacement Over the Business Cycle and Its Sources: Evidence from Germany," *Working Paper*, 2018.

Schultz, Theodore W., "The Value of the Ability to Deal with Disequilibria," *Journal of Economic Literature*, 1975, *13* (3), 827–846.

Seim, David, "On the incidence and effects of job displacement: Evidence from Sweden," *Labour Economics*, 2019, *57* (April), 131–145.

Shah, Manisha and Bryce Millett Steinberg, "Drought of Opportunities: Contemporaneous and Long-Term Impacts of Rainfall Shocks on Human Capital," *Journal of Political Economy*, 2017, *125* (2), 527–561.

Shastry, Gauri Kartini, "Human capital response to globalization education and information technology in India," *Journal of Human Resources*, 2012, *47* (2), 287–330.

Small, Michael H., "Assessing manufacturing performance: an advanced manufacturing technology portfolio perspective," *Industrial Management & Data Systems*, 1999, *99* (6), 266–278.

Smith, Adam, *An Inquiry into the Wealth of Nations*, London: Strahan and Cadell, 1776.

Sohal, Amrik S., "Assessing AMT implementations: an empirical field study," *Technovation*, 1996, *16* (8), 377–444.

Solow, Robert M., "A Contribution to the Theory of Economic Growth," *The Quarterly Journal of Economics*, 1956, *70* (1), 65–94.

Spitz-Oener, Alexandra, "Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure," *Journal of Labor Economics*, 2006, *24* (2), 235–270.

Stapleton, Katherine and Michael Webb, "Automation, Trade and Multinational Activity: Micro Evidence from Spain," *Working Paper*, 2020.

Stevens, Ann Huff and Jessamyn Schaller, "Short-run effects of parental job loss on children's academic achievement," *Economics of Education Review*, 2011, *30* (2), 289–299.

Stokey, Nancy L., Robert E. Lucas, and Edward C. Prescott, *Recursive Methods in Economic Dynamics*, Harvard University Press, 1989.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Sutton, John, *Technology and Market Structure: Theory and History*, Cambridge: MIT Press, 1998.

Takalo, Tuomas, Tanja Tanayama, and Otto Toivanen, "Estimating the Benefits of Targeted R&D Subsidies," *The Review of Economics and Statistics*, 2013, *95* (1), 255–272.

Taylor, Frederick Winslow, *The Principles and Methods of Scientific Management*, New York and London: Harper & Brothers, 1911.

Tinbergen, Jan, *Income Difference: Recent Research*, North-Holland Publishing Company, 1975.

Tolbert, Charles M and Molly Sizer, "US Commuting Zones and Labor Market Areas: A 1990 Update," *U.S. Department of Agriculture Economic Research Service Staff Paper*, 1996.

Utterback, James M. and William J. Abernathy, "A dynamic model of process and product innovation," *Omega*, 1975, *3* (6), 639–656.

Vainik, Uku, Alain Dagher, Anu Realo, Lucía Colodro-Conde, Erik Lykke Mortensen, Kerry Jang, Ando Juko, Christian Kandler, Thorkild IA Sørensen, and René Mõttus, "Personality-obesity associations are driven by narrow traits: A meta-analysis," *Obesity Reviews*, 2019, *20* (8), 1121–1131.

Vance, J. D., *Hillbilly Elegy: A Memoir of a Family and Culture in Crisis*, Harper Press, 2016.

Verhoogen, Eric A, "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *The Quarterly Journal of Economics*, 2008, *123* (2), 489–530.

von Wachter, Till and Elizabeth Weber Handwerker, "Variation in the Cost of Job Loss by Worker Skill: Evidence Using Matched Data from California, 1991–2001," *Working Paper*, 2009.

Weaver, Andrew and Paul Osterman, "Skill Demands and Mismatch in U.S. Manufacturing," *ILR Review*, 2017, *70* (2), 275–307.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," *Working Paper*, 2020.

Weinberger, Catherine J, "The increasing complementarity between cognitive and social skills," *Review of Economics and Statistics*, 2014, *96* (4), 849–861.

Welch, Finis, "Education in Production," *Journal of Political Economy*, 1970, *78* (1), 35–59.

Whyte, William Foote, *Street Corner Society*, University of Chicago Press, 1943.

Willis, Paul, *Learning to Labour: How Working Class Kids Get Working Class Jobs*, Columbia University Press, 1977.

Wilson, William Julius, *When Work Disappears*, New York: Alfred A. Knopf, 1996.

Xiang, Chong, "New Goods and the Relative Demand for Skilled Labor," *The Review of Economics and Statistics*, 2005, *87* (2), 285–298.

Yagan, Daniel, "Employment Hysteresis from the Great Recession," *Working Paper*, 2017.

Zhang, Tong, Fred Damerau, and David Johnson, "Text Chunking based on a Generalization of Winnow," *Journal of Machine Learning Research*, 2002, *2* (Mar), 615–637.