Year : 2022

Three Essays in Labor Economics

Colella Fabrizio

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

THREE ESSAYS IN LABOR ECONOMICS

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Docteur ès Sciences Économiques,
mention « Économie politique »

par

Fabrizio COLELLA

Directeur de thèse
Prof. Rafael Lalive

Jury

Prof. Christian Zehnder, président
Prof. Mathias Thoenig, expert interne
Prof. David Card, expert externe
Prof. David Dorn, expert externe

LAUSANNE
2022

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

THREE ESSAYS IN LABOR ECONOMICS

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Docteur ès Sciences Économiques,
mention « Économie politique »

par

Fabrizio COLELLA

Directeur de thèse
Prof. Rafael Lalive

Jury

Prof. Christian Zehnder, président
Prof. Mathias Thoenig, expert interne
Prof. David Card, expert externe
Prof. David Dorn, expert externe

LAUSANNE
2022

UNIL | Université de Lausanne
HEC Lausanne
Le Décanat
Bâtiment Internef
CH-1015 Lausanne

IMPRIMATUR

Sans se prononcer sur les opinions de l'auteur, la Faculté des Hautes Etudes Commerciales de l'Université de Lausanne autorise l'impression de la thèse de Monsieur Fabrizio COLELLA, titulaire d'un bachelor en économie et finances et d'un master en sciences économiques et sociales de l'Université Luigi Bocconi de Milan, et d'un master en analyse économique de l'Université Carlos III de Madrid, en vue de l'obtention du grade de docteur ès Sciences économiques, mention « économie politique ».

La thèse est intitulée :

THREE ESSAYS IN LABOR ECONOMICS

Lausanne, le 9 mars 2022

La Doyenne

Marianne SCHMID MAST

Members of the doctoral committee

Professor Rafael Lalive
Professor of Economics, HEC, University of Lausanne
Thesis supervisor

Professor Mathias Thoenig
Professor of Economics, HEC, University of Lausanne
Internal member of the doctoral committee

Professor David Card
Class of 1950 Professor of Economics, University of California, Berkeley
External member of the doctoral committee

Professor David Dorn
UBS Professor of Globalization and Labor Markets, University of Zurich
External member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics
Subject area Political Economy

I hereby certify that I have examined the doctoral thesis of

Fabrizio COLELLA

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: R. Lalive Date: 16.2.2022

Prof. Rafael LALIVE
Thesis supervisor

University of Lausanne
Faculty of Business and Economics

PhD in Economics
Subject area Political Economy

I hereby certify that I have examined the doctoral thesis of

Fabrizio COLELLA

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: _____________________________ Date: 27/02/2022

Prof. Mathias THOENIG
Internal member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics
Subject area Political Economy

I hereby certify that I have examined the doctoral thesis of

Fabrizio COLELLA

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ____________ Date: 16/2/2022

Prof. David CARD
External member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics
Subject area Political Economy

I hereby certify that I have examined the doctoral thesis of

Fabrizio COLELLA

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: D. Dorn Date: 17.02.2022

Prof. David DORN
External member of the doctoral committee

Three Essays in Labor Economics

by

Fabrizio Colella

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to my advisor Rafael Lalive for his generous guidance. He mentored me since the first day I entered the Department and has been my reference point during all my academic career. He constantly encouraged me and taught me how to deal with every circumstance. I owe him much of my professional knowledge and most of the progresses I made during my PhD studies. Rafael is the kind of advisor that every PhD student wants to have.

I am extremely grateful to David Card for his unconditional support and for giving me the priceless opportunity to visit UC Berkeley. I learnt a lot from every single interaction with him, he is a constant source of inspiration as researcher and person.

I am truly indebted to Mathias Thoenig, who played a key role in shaping my training. My approach toward the profession is profoundly influenced by his hardworking and precise attitude.

I would like to thank also David Dorn for his invaluable suggestions on the chapters of this dissertation and for driving my research interests with his inspiring papers.

This dissertation extensively benefited from comments and suggestions from David Autor, Hélène Benghalem, Marius Brülhart, Johannes Buggle, Patricio Dalton, Rustamdjan Hakimov, Peter Kuhn, Omar Rachedi, Tiziano Rotesi, Enrico Rubolino, Seyhun O. Sakalli, Christofer Schroeder, Camille Terrier, Andrea Weber, and Josef Zweimüller. I would like to sincerely thank them all for the precious help.

Finally, I would like to thank also the faculty, the postdoctoral fellows and the PhD students of the Department of Economics of the University of Lausanne as well as the visiting and the postdoctoral researchers at Center for Labor Economics and at the Institute for Research on Labor and Employment of the University of California, Berkeley I interacted with. These people kindly accompanied me during this journey providing a constant source of stimulating conversations and joy.

Abstract

The three chapters of this thesis contribute to the understanding of how economic shocks and discriminatory behaviors of recruiters and stakeholders impact the allocation of workers in different jobs as well as their performance at work.

Chapter one examines the extent to which changes in international market prices lead to shifts in firms' skill requirements through trade. I exploit an exceptional change in international prices generated by a decision of the Swiss National Bank, which caused a 15% increase in the value of the local currency. I study how this sharp change in trade conditions affected skill requirements in Switzerland using novel data on trade and labor demand. I find that, in the two years after the shock, firms with a workforce more exposed to offshorability and automation increased imports and posted more job advertisements for highly skilled workers.

Chapter two, coauthored with David Card and Rafael Lalive, studies how forbidding the use of stated gender preferences in job advertisements impacts workplace gender diversity. We take advantage of a campaign in Austria that almost eliminated all job openings specifying a gender preference. We show that prior to the campaign, most stated preferences were concordant with the firm's existing gender composition and they were very likely to be filled by someone of that gender. We find that the elimination of gender preferences increased the diversity of hiring workplaces and did not deteriorate job match efficiency.

Chapter three investigates the effect of supporters on the performance of soccer players by skin color. Due to the COVID-19 restrictions, one third of the games of the highest Italian soccer league 2019/2020 season were played in closed stadiums. By using objective player performance data and an automated skin color recognition algorithm, I identify a significant increase in the performance of non-white players, relative to white players, when supporters are banned from the stadium.

Contents

List of Figures

List of Tables

INTRODUCTION

The economy constantly evolves: every day new and more sophisticated goods are produced and the way in which they are exchanged by economic operators continuously improves. Innovation and economic shocks encompass important changes to societies and have consequences in terms of labor market outcomes. Most of labor market adjustments transit through employers' decisions, since they play a crucial role in shaping the allocation of workers and their wages.[1] Firms are the main actors that model labor demand and have some power in setting workers' salaries, as pointed by Card (2022) during the latest Annual Meeting of the American Economic Association.

When structural changes generate new incentives for firm, such as new selling opportunities, or the possibility to adopt more efficient technology or to acquire cheaper inputs from abroad, workers are affected. In the first chapter of this thesis I study how a monetary shock leads to shifts in firms' skill requirements through trade. On January 15, 2015 the Swiss National Bank unexpectedly abandoned the exchange rate floor with the Euro, causing a 15%increase in the value of the Swiss franc, which remained relatively stable in the subsequent years. This unforeseen appreciation immediately impacted the relative price of trade, increasing the incentive to import, while simultaneously reducing expected profits for firms exposed to foreign competition. I merge trade data containing information on each import or export transaction made by Swiss firms with firm-specific job postings data. I construct a new firm-specific pre shock measure of labor force substitutability and I find that exposed firms increased imports during the two years after the shock and that this increase in imports translates into a reduction in the demand for routine-intensive tasks and manufacturing skills. For these firms, a 10% increase in monthly imports translates into a 2.1% reduction in the

[1] For example, it has been extensively documented how the adoption of computers by firms at the end of the last century generated important shifts in the employment and increased the wage gap between certain occupations (Autor and Dorn, 2013).

routine intensity of the tasks associated with their labor demand.

While firms needs are the driving force guiding their hiring, recruiters' preferences and stereotypical believes play an important role as well. Policymakers strive to prohibit behaviors governing gender roles in order to give equal opportunities to all workers. However, women and men still tend to work in different workplaces. In the second chapter, coauthored with David Card and Rafael Lalive, we investigate how discriminatory employer preferences impact workplace gender segregation and hiring efficiency. In spring 2005, Austria launched a campaign to inform employers and newspapers that gender preferences in job advertisements were illegal, reducing the share of openings on the nation's largest job-board specified a preferred gender from 40% to 5%. We merge data on filled vacancies to linked employer-employee data to study how the elimination of gender preferences affected hiring and job outcomes. Before the campaign most stated preferences were concordant with the firm's existing gender composition. We use pre-campaign vacancies to predict the probabilities of specifying preferences for females, males, or neither gender. We then conduct event studies of the effect of the campaign on the predicted preference groups. We find that the elimination of gender preferences led to a rise in the fraction of women hired for jobs that were likely to be targeted to men (and vice versa), increasing the diversity of hiring workplaces. Vacancy filling times, wages, and job durations were largely unaffected by the campaign for most of the vacancies.

Discriminatory behaviors, driven by individual preferences, do not impact only how workers allocate into different jobs, but influence also workers' performance. This phenomenon is common in sport events, where discrimination manifests through racist harassment of supporters toward athletes. In the third chapter of this thesis I study how racist behaviors from supporters affect the performance of soccer players in Italy. By forcing teams to play games without fans, the COVID-19 pandemic created a compelling natural experiment. I generate an automated skin color recognition algorithm to classify players in white and non-white. Next, I compile individual performance scores generated by an algorithm used for fantasy-sports competitions. I show that, on average, white players record a lower score without fans than they did in packed stadiums, while non-white footballers' improved their performances. The effect does not differ between home and away games, and players playing in top versus minor teams, while weaker players are impacted more than others.

CHAPTER 1

THE EFFECT OF TRADE ON SKILL REQUIREMENTS: EVIDENCE FROM JOB POSTINGS

1.1 INTRODUCTION

Labor markets are evolving rapidly, encompassing dramatic changes in the tasks that individuals perform at work along with increases in wage differentials across jobs (Autor and Dorn, 2013; Deming and Noray, 2020).[1] A growing line of research documents how international trade affects the labor demand and generates important changes in employment and wages (Hummels et al., 2014; Dorn and Levell, 2021). There is indeed a common consensus that the opening of the domestic market to new countries has led to higher unemployment in import-competing sectors in both the United States (Autor et al., 2013; Autor et al., 2014) and Europe (Goos et al., 2014; Dauth et al., 2014).

Exchange rate movements have a strong bearing on export and import prices, impacting firm decisions towards trade and consequently employment, wages (Campa and Goldberg, 2001), and skill requirements (Kaiser and Siegenthaler, 2016). When a country experiences a currency appreciation, domestic firms become less competitive

[1] An extensive literature documents how part of the shift in the labor demand toward highly skilled workers is related to the emergence of new machines (Krusell et al., 2000; Autor et al., 2003). This literature coined the concept of skill biased technological change (SBTC), defined as a change in the production function toward a new technology that complements skilled workers. See Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018a) for an overview of this literature.

in both domestic and foreign markets. On the one hand, the consequent contraction of the sector's market share generates detrimental effects on its employment in exposed sectors. On the other hand, the adoption of foreign inputs and technology becomes cheaper, creating incentives for a reorganization of the production chain that incorporates more imported inputs. Currency shocks provide a suitable setting to study these dynamics (e.g. Verhoogen, 2008). While the pass-through effect of a currency appreciation, or depreciation, on prices has been largely investigated, little is known on how exogenous exchange rate shocks impact firms' trading decisions and the composition of their labor demand.

In this paper, I take advantage of a sudden and unexpected appreciation of the Swiss franc to study how changes in trade opportunities generated by a price variation influence skill requirements. for Swiss manufacturing firms. In September 2011, the Swiss National Bank introduced an exchange rate floor of 1.20 Swiss francs per Euro, and the rate was almost always binding in the subsequent years. On January 15, 2015 this floor was unexpectedly abandoned and the Swiss franc immediately appreciated by more than 20%, and then stabilized around a 15% higher value for the following years. In one day, firms that were exclusively exporters saw their expected revenues from foreign sales potentially cut by 15%, while those that were exclusively importers had a sudden expected 15% cost saving on their imported products. By lowering the price of imported inputs, the shock has generated an incentive for Swiss firms to offshore part of their production process or to import cheaper capital from abroad. Both adjustment strategies could increase the demand for high skilled labor, by either substituting workers involved in offshorable activities with foreign inputs or by substituting those that are at risk of automation.

I use transaction-level custom's data containing all import and export movements at the firm level to track the foreign trade of Swiss firms around the shock. I merge this data with firm-specific online job advertisements to record changes in the labor demand of different firms.[2] With this, I develop a new measure of exposure to the shock, combining the offshorability and routine intensity of the pre-shock labor force and classify firms according to it. I then estimate a series of difference-in-differences

[2]Online job postings are shown to be useful for studying variations in employer skill demands within labor markets and occupations (Modestino et al., 2016; Deming and Kahn, 2017; Atalay et al., 2018; Azar et al., 2018 and Hershbein and Kahn, 2018).

(DiD) models to calculate the impact of the shock on firms whose workers are exposed to substitutability versus those that are not. The outcomes are different indicators of trade and labor demand composition. To dig deeper into the effect of trade on skill requirements, I leverage exogenous changes in imports, triggered by the appreciation shock for exposed firms. Specifically, I estimate a two stage least square model instrumenting imports with an interaction term of the shock and the exposure measure.

I find a positive and significant correlation between trade, labor demand and indicators of skill requirements in the pre-shock period. Comparing pre and post shock levels, I show that a one-standard-deviation increase in exposure to substitutability translates into an increase in monthly imports by 8.3% and increases the demand for high skilled workers. For each additional log monthly import, a firm reduces the routine intensity associated with its labor demand by 0.137 percentage points (2/3 of standard deviation) and the share of jobs requiring manufacturing skills by 0.107 percentage points (1/2 of standard deviation). There is no statistically significant effect of the shock on changes in export by exposure. To the best of my knowledge, this is the first paper that combines job posting data with custom's data on trade at the firm level and that provides evidence on the effect of currency appreciation on detailed skill requirements.

This paper contributes to several strands of the literature. Primarily, it relates to the literature on the effect of trade on labor market outcomes. This line of research documents long term changes in employment, occupational differences and wages induced by trade and import competition. Autor et al. (2013) and Autor et al. (2014) provide a detailed description of the evolution of the US labor market after China joined the World Trade Organization (WTO). The ease of access to foreign markets generated higher unemployment, lower labor force participation and a reduction in lifetime earnings for workers employed in import-competing manufacturing industries. Goos et al. (2014) reveal evidence of similar patterns in Europe, while Dauth et al. (2014) demonstrate that employment gains for exported oriented sectors outweigh the job losses from competition with foreign firms.

A channel through which trade affects skill requirements is innovation. Building on the seminal paper by Acemoglu (1998), Hanlon (2015) provides evidence that changes in the relative supply of inputs drives the direction of technological process. Bustos (2011) documents an increase in technology upgrading for firms facing tariff reductions in Argentina due to a free agreement with Brazil. Bloom et al. (2016) emphasize how exposure to import competition from China generates innovation in Europe and shifts employment towards more technologically advanced firms. Thoenig and Verdier (2003) examine skill-biased effects of trade through technological change. In their model, firms respond to foreign competition by shifting innovations towards skilled labor intensive technologies, contributing to an increase in occupational wage inequality. Tanaka (2020) investigates the rapid opening of Myanmar to foreign trade after 2011 and finds that exporting has large positive impacts on sales, employment, working conditions, and wages. Another channel by which trade generates an impact on labor market is offshoring. Analyzing Danish firms and workers, Hummels et al. (2014) estimate that offshoring increases wages for highly skilled workers and generates wage losses for routine-based jobs. In line with this, Hijzen et al. (2005) show that international outsourcing has had a strong negative impact on the demand for unskilled labor in the United Kingdom, while Becker et al. (2013) associate offshoring with a shift in labor demand towards more non-routine and interactive tasks, and highly educated workers. This paper integrates these two lines of research by causally identifying skill-biased effects of changes in trade opportunities for firms exposed to automation or offshorability.

Currency fluctuations are often used to identify the relationship between trade and firm decisions. Bastos et al. (2018) exploit exchange-rate movements in Portugal as a determinant for firm decisions about investment location. They find that exporting to richer countries leads firms to pay higher prices for inputs. Alvarez and Lopez (2015) investigate the relationship between the real exchange rate and the acquisition of foreign technology in Chile. They document that a real depreciation of the local currency increases foreign technology acquisition, but only among exporting firms. Chen (2017) uses country-level data for 49 countries and shows that an undervaluation of the real exchange rate decreases R&D expenditure. He primarily attributes this effect to the increasing costs of importing machinery and other inputs due to the currency depreciation. Kaiser and Siegenthaler (2016) use a panel of Swiss manufacturers covering the period 1998-2012 to show that an upward fluctuation of the local currency increases high-educated and reduces low-educated employment. While highlighting important employment effects of exchange rates on labor demand, all these papers rely on currency fluctuations rather than a cleaner discontinuity in prices. One exception is the paper by Verhoogen (2008), who analyzes the labor market effects of the 1994 peso crisis shock. He finds that initially more-productive plants increased the export share of sales, white-collar wages, blue-collar wages, and the relative wage of white-collar workers relative to initially less-productive plants. I complement this literature by analyzing a sharp variation in trade opportunities provided by the Swiss franc shock. In contrast to Verhoogen (2008), I exploit a currency shock occurring in an otherwise stable period for the Swiss economy. Further, I study a sudden currency appreciation rather than a currency crisis, which, in theory, is supposed to have an opposite effect in terms of firms' incentives compared to the ones documented through the peso's crisis. This paper differs from this branch of literature also with regard to outcomes: the analysis of job postings, rather than employment, allows for a detailed description of the labor demand by firms across the skills distribution.

This paper also relates to the literature on the evolution of skill requirements and its determinants. The acquisition of new technologies is shown to have substantial effects in shifting labor demand. For this reason, this literature directly relates to the concept of skill biased technological change (SBTC)and to the *computerization* phenomenon. Krusell et al. (2000) show that the canonical SBTC framework can explain a large part of the skill premium. However, Card and DiNardo (2002) provide evidence that SBTC fails in explaining the existing wage differentials between different groups of workers. Autor et al. (2003) argue that computer capital is a substitute for routine, algorithmic, low and middle skill tasks, while it is a complement for cognitive, problem solving, and communication tasks. Autor and Dorn (2013) demonstrate how employment and wage polarization in the United States labor market partially stems from the adoption of information technology and the consequent falling cost of automation. Goos and Manning (2007) find similar results for Britain. I add to this literature by showing an additional channel through which the demand for skills evolves in a small open economy: the change in international market prices due to an unforeseen currency appreciation.

This paper also contributes to the growing literature that exploits information contained in the text of job advertisements to study labor market concentration (Azar et al. (2018)) and to better understand labor demand dynamics. Most of this research was made possible through the data on US job postings provided by Burning Glass Technology. Online job postings are shown to be useful for studying variations in employer skill demands across labor markets and occupations as they contain a detailed description of the requirements and are observed at high frequency and at a refined geographical level with respect to survey or administrative data. The seminal paper by Deming and Kahn (2017) shows the existence of substantial variation in skill requirements within occupations and that skill requirements positively correlate with wages. Hershbein and Kahn (2018) analyze the change in skill demand induced by the Great Recession, showing that US metropolitan statistical areas that suffered a higher unemployment rate during the crisis had a relatively greater increase in skill requirements. Modestino et al. (2016) document a negative relation between employer skill requirements and the business cycle. I exploit, for the first time, a new database on job postings from Switzerland that is comparable to the Burning Glass Technology database in terms of contents. I follow the literature in the use of skill content extracted from job advertisements through text analysis, and complement it by showing how skill requirements in job ads are affected by a currency appreciation.

Finally, this paper adds to the literature exploiting the extraordinary variation in real prices provided by the Swiss exchange rate shock. Several papers study the exchange rate pass-through on prices (Auer et al., 2021; Auer et al., 2019; Kaufmann and Renkin, 2019; Bonadio et al., 2019). Efing et al. (2015) investigate the impact of the Swiss franc shock on investments by publicly listed Swiss firms. They find that firms with large currency risk exposure decreased their real investments by 8.1% half a year after the abolition of the exchange rate floor. Kaiser et al. (2018) employ a DiD approach by comparing the evolution of the investment of firms with different net exposure before and in the two years following the Swiss franc shock. They find that firms with positive net exposure reduced gross fixed capital investment by roughly 15% in 2015 and by 12% in 2016 relative to negatively exposed firms. Exposed firms reduced investment in machinery and equipment, construction and R&D in 2015 and 2016. They also find evidence suggesting that the franc shock appears to have induced exposed firms to renew their machinery and equipment. The closest paper to this one is a policy report written by Kaufmann and Renkin (2017) for the Swiss State Secretariat for Economic Affairs. They employ a DiD approach comparing Swiss manufacturing firms with a control group of similar Austrian firms before and after the Swiss franc shock. They find an average 4% decrease in employment for Swiss manufacturing firms in the two years after the appreciation. They also find a reduction in job vacancies and show how the decline in vacancy postings explains most of the variation in employment. However, they do not provide any evidence on potential skill-biased effects of the shock on labor demand. I integrate this literature by investigating the effect of the Swiss franc shock on international trade and skill requirements, using detailed information on the demand for skills and firms' products import and export, before and after the shock.

The remainder of this paper proceeds as follows. Section 1.2 provides background information on the Swiss franc shock and its implications. Section 1.3 describes the data, while Section 1.4 investigates the correlation between trade and labor demand and their evolution. Section 1.5 provides details on the empirical strategy and presents evidence that the trade induced by the exchange rate shock has a skill-biased impact on labor demand. Section 1.6 concludes.

1.2 Background

1.2.1 The Swiss Manufacturing Sector

Switzerland's manufacturing sector is a central pillar of the Swiss economy. According to the World Bank (2020), it generates about 18% of the Swiss GDP and it is the second largest employer in the country. Swiss manufacturing specializes primarily in high-tech and knowledge-based production. Major products include machinery and equipment, chemical-pharmaceutical products as well as scientific and precision instruments, such as luxury watches and hearing aids.[3] The largest industry is mechanical engineering, electrical engineering and metalworking (MEM), which makes up almost half of the manufacturing sector. With around 320,000 employees, it is the biggest industrial employer in Switzerland (SWISSMEM, 2019). This sector has increasingly become more dedicated to dynamic technology fields such as sensor technology, photonics, robotics, additive manufacturing, and industrial IT. This industry is highly export oriented, and the main destination is the European Union, accounting for 60% of total exports, followed by the United States at 14%.

Swiss products and services have a good reputation worldwide. Customers associate Swiss made products with reliability, quality, longevity and technological superi-

[3] In comparison, the largest manufacturing industries in the US for instance include petroleum, steel, automobiles, aerospace as well as food processing, consumer goods and electronics.

ority. The “Swiss Made" quality indicator in particular benefits manufacturing companies in the business to business sector. To earn the “Swiss Made" label, a company's industrial product must incur at least 60% of its manufacturing costs (including R&D, material and production costs including costs for quality assurance and certification) in Switzerland. In addition, the main production stage must take place in Switzerland (Switzerland Global Enterprise, 2020).

Switzerland's workforce is highly skilled and the manufacturing sector has capitalized well on this. The large stock of high quality labor not only provides a steady stream of qualified workers but also a consistently high added-value to the economy through specialization. The main reason behind this extremely specialized skilled technical workforce is the Swiss vocational education system. Specifically, the vocational training system in Switzerland is oriented toward the labor market and incorporates both theory and practice (Switzerland Global Enterprise, 2020). Consequently, thousands of young, well-trained individuals enter the labor market every year, particularly in the manufacturing industry. While some workers are highly specialized and hard to replace, the liberal labor regulations with relatively weak job protection allows firms to employ and dismiss staff at short notice without incurring incidental wage costs.

Historically, the Swiss manufacturing sector has been more resilient to foreign competition than other countries. High-margin industries such as pharmaceuticals and watchmaking have been able to focus on value creation and export growth in recent past years, and have increased productivity and employment rather than reducing the size of the workforce. The continuous innovation of Swiss manufacturing companies makes them less vulnerable to competition at the output level and prone to accomodate automation.

1.2.2 The Swiss Franc Shock

The appreciation of the Swiss Franc was the result of a sudden monetary policy decision by the Swiss National Bank (SNB). After the global financial crisis, the Swiss franc (CHF) saw a strong appreciation with respect to the euro (EUR) and the US dollar. This process accelerated with the following euro crisis generating a drop in the CHF/EUR rate by more than 30% in 2 years. To curb this excessive appreciation, the SNB implemented a non-conventional monetary policy measure, by announcing on September

6th, 2011 a CHF/EUR exchange rate floor of 1.20. The SNB stated that it was ready to buy an unlimited amount of foreign currency to maintain the floor, if necessary. As a consequence, the CHF/EUR stabilized in the 1.20-1.24 window for the subsequent years, often binding at 1.20. On January 15, 2015, the SNB ended the exchange rate floor policy announcing that it would no longer artificially keep the Swiss franc low. After this announcement the Swiss franc appreciated against the Euro and other currencies. The CHF/EUR exchange rate instantaneously dropped to 0.98 and after a high volume of transactions it remained around 1.04 Swiss francs per Euro in the following months. Figure 1.1a plots the EUR/CHF exchange rate in the 3 years before and after the shock, showing that the ratio fluctuated between 1.05 and 1.10 even in subsequent years without ever going back to the pre-shock levels.

The Swiss franc shock provides an ideal and unique setting to study trade and its effects on the labor market for at least three reasons. First, the SNB decision to stop artificially maintaining a floor level was unanticipated and perceived as permanent. Figure 1.1b plots the EUR/CHF around the shock, together with stock market operators' 3-month predictions of the exchange rate highlighting that markets did not expect the end of the floor by the SNB. Data on expectations are taken from Mirkov et al. (2016), who analyzed financial derivatives on the exchange rate and transactions on the Forex liquidity market and show no significant shift in market expectations. Figure 1.1b shows also that expectations adjusted quickly after the shock, signaling that operators perceived the change as persistent. The KOF Consensus Forecast surveys a panel of 20 economists quarterly, asking them to forecast Swiss franc value with respect ot the euro for the following 12 months. One month before the shock, the average forecast measure predicted that the rate would stay at about 1.2 Swiss francs per euro in the following year, with a narrowed confidence interval. Also in this case expectation adjusted immediately, as shown in figure 1.A1a, taken from Kaufmann and Renkin (2019). Further, figure 1.A1b in appendix shows that general interest toward the topic was also quite stable before the SNB announcement.

Second, this unexpected monetary policy change occurred in a period of relatively stable economic circumstances and after a three year period of stability. This allows for a clear before-after comparison and allows to discern the effect of trade from other confounders.

Third, contrary to many currency crisis shocks, the result of the January 15, 2015

Figure 1.1: The Swiss Franc Shock

(a) EUR/CHF Exchange rate around the shock

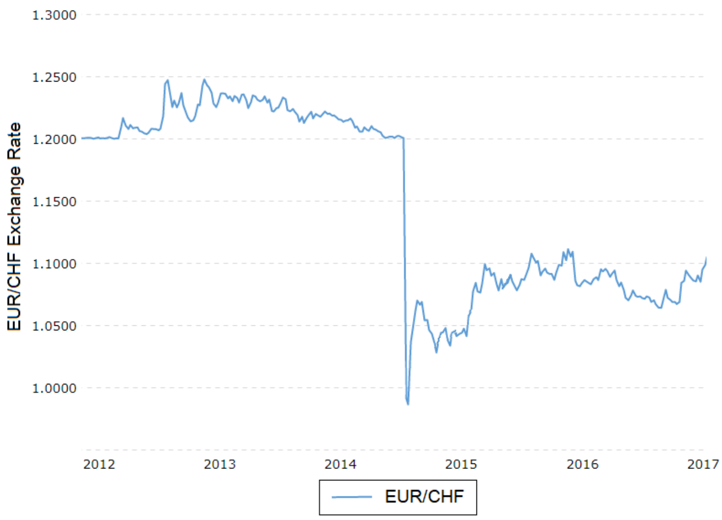

(b) EUR/CHF Exchange rate around the shock

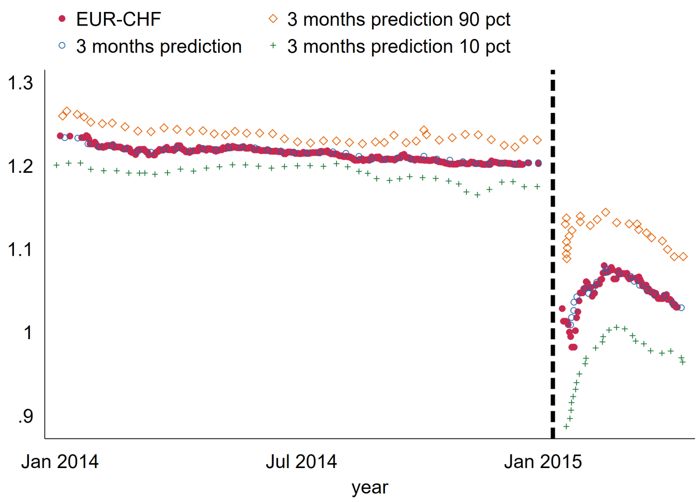

Notes: The figure in panel (a) plots the daily EUR/CHF exchange rate in the three years before and the three years after the shock. The figure in panel (b) plots the EUR/CHF around the shock and the 3-month predictions of the exchange rate together with its confidence interval. Expectations are constructed using the indicators built by Mirkov et al. (2016) based on stock market prices of derivatives.

SNB decision was a sudden appreciation of the local currency, making imports cheaper and exports more expensive. This is crucial for the study of the labor market and the skill requirements, as it directly impacts the production function of firms creating incentives for offshoring and automation. The shock is particularly relevant since the Eurozone is the most important market for Swiss firms, accounting for more than 60% of the Swiss international manufacturing trade.

1.2.3 Implications

Motivated by the low cost of labor in foreign countries and the need to compete with emerging economies, offshoring has been a common strategy for manufacturing companies based in developed countries over the last 30 years. Firms disaggregate their value chain, outsourcing some tasks to external providers located abroad. The external provider produces inputs that were previously produced in-house and sells them to the outsourcing company. The effect of this restructuring of the firm production function on its workforce is likely to be skill-biased (Hanlon, 2015). Offshoring in manufacturing usually concerns standardized products, requiring mainly non-cognitive and manual skills. Moreover, to offshore the firm needs to negotiate with foreign firms and to organize the internal structure, which requires managerial skills (Hijzen et al., 2005; Biscourp and Kramarz, 2007; Becker et al., 2013; Hummels et al., 2014). This implies that imported inputs substitute for unskilled labor while complementing skilled workers, meaning that offshoring should upward bias firm relative labor demand in terms of skills.

An appreciation of the exchange rate lowers the relative price of imported goods making outsourcing more convenient with respect to the other domestic production factors. A survey conducted in August 2015 by Swissmem, the leading association in the Swiss industry, shows how Swiss firms consider offshoring as reaction to the currency shock. Almost half (46%) of the industrial companies surveyed have decided more or less firmly to relocate part of their activities to the European Union (EU) and to make investments outside Switzerland, while 18% of companies were considering moving their entire production outside Switzerland.[4] As a consequence, the relative

[4]The results of the survey were published online on August 20, 2015, by the newspaper *Tribune de Genève*: *www.tdg.ch/economie/entreprises/delocalisation-tente-industriel-deux.*

demand for highly skilled workers should increase after a positive currency shock (Kaiser and Siegenthaler, 2016; Burstein et al., 2013).[5]

In addition, most of the the capital used by Swiss firms is imported from abroad. Therefore, by making foreign capital less expensive, the currency appreciation might affect skill requirements through a second channel. The appreciation could trigger the import of new technology, which is shown to be complementary to highly skilled labor (Acemoglu, 1998; Krusell et al., 2000).

To illustrate how the exchange rate shock might impact labor demand, I augment the model in Krusell et al. (2000) introducing offshored labor as a close substitute for low skill labor, as in Hummels et al. (2014).[6] I assume that a representative firm can produce a good y using five different production factors. Three factors that are acquired domestically, capital structure k_d, high-skill labor h and low-skill labor l, and two factors imported from abroad and are offshored labor o and imported capital equipment k_i.[7] The production function is the following.

$$y = {k_d}^{\alpha} \left\{ \mu \left[\beta o^{\theta} + (1-\beta) l^{\theta} \right]^{\frac{\sigma}{\theta}} + (1-\mu) \left[\lambda {k_i}^{\rho} + (1-\lambda) h^{\rho} \right]^{\frac{\sigma}{\rho}} \right\}^{\frac{1-\alpha}{\sigma}} \tag{1.1}$$

The parameters μ is the weight associated to low skill labor, β is the weight associated to offshored labor and λ is the weight associated to foreign capital. θ, σ, and ρ are parameters that govern the degree of substitutability or complementarity of the five factors of production.

The production function depends not only domestic capital but also on two other

[5] As a small open economy, Switzerland is subject to international price spillovers which depend on the underlying shocks abroad (Baurle et al., 2017). However, the Swiss export market is inherently special, as exports are relatively inelastic to exchange rate and price fluctuations (Auer et al., 2021), while imports are not. Hanslin Grossmann et al. (2016) conclude that not only exports were less sensitive to exchange rate fluctuations, but that the underlying export sectors shifted towards a structure that is more independent of business cycles.

[6] An alternative would be following the model of Acemoglu and Restrepo (2018b) and include offshored labor. However, their setting in which technology adoption destroys some jobs and creates new tasks to be performed by workers would suit better for an investigation on the new skills requirements. In my analysis I focus on some indicators of skill requirements and study how their demand changes around the shock. However, I intend to disentangle the intensive and the extensive margins effects of the shock on labor demand in future iterations.

[7] Offshored labor refers to foreign labor force or imported intermediate inputs. When offshoring concerns the production of goods, rather than services, the two concepts are the same, as usually firms imports the inputs that are produced by the foreign workers. I could also include domestic capital equipment in the production function as a close substitute for the imported capital equipment, but this would not change the implications. In the rest of the paper I use interchangeably capital structure and domestic capital as well as capital equipment and imported capital.

components, one is a CES composite of low-skill and offshored labor, $X = [\beta o^{\theta} + (1-\beta) l^{\theta}]^{\frac{1}{\theta}}$, and the other component is a CES composite of imported capital and high-skill labor, $Z = \left[\lambda k_i{}^{\rho} + (1-\lambda) h^{\rho}\right]^{\frac{1}{\rho}}$. What is key of this production function is that it is flexible enough to have different elasticity of substitutions between and across inputs. As in Krusell et al. (2000), the production function has constant returns to scale nested CES in all the inputs and it has a unitary elasticity of substitution between domestic capital k_d and a nested CES function of the two other CES composite inputs. The elasticity of substitution between low-skill and offshored labor is given by $1/(1-\theta)$, the one between high-skill labor and imported capital is $1/(1-\rho)$ and the one across the two bundles X and Z is $1/(1-\sigma)$.

The function restricts the elasticity of substitution between the composite good X and high-skill labor to be the same as that between X and imported capital.[8] Analogously, it restricts the the elasticity o substitution between the composite good Z and low-skill labor to be be equal to that between Z and offshored labor. Three other restrictions drives the complementarity and substitutability among inputs: $1/(1-\theta) > 1$ implies close substitutability between low-skill and offshored labor, $1/(1-\sigma) < 1$ implies complementarity between the low-skill/offshored labor bundle X and the capital/high-skill bundle X, $1/(1-\rho) < 1/(1-\sigma)$ imposes the complementarity between imported capital and high skill to be stronger than the one between imported capital and the boundle X.

The profit function associated with equation 1.1 is the following:

$$\Pi = y_d\, p^D + y_x\, e p^X - h\,\omega_h - l\,\omega_l - k_d\, p_d - o\, e\omega_o - k_i e p_i \tag{1.2}$$

The prices of each unit of high-skill and low-skill labor are ω_h and ω_l, while the rental price of domestic capital is p_d. The prices of the two imported factors of production, in the foreign currency, are w_o for outsources labor, and p_i for imported capital. In terms of the domestic currency, the prices are $e\omega_o$ and ep_i, respectively; e it the exchange rate franc per euro. The final produced good can be sold domestically at the price p^D and abroad at the price p^X.

The exchange rate shock can be interpreted in the model as a drop in e, which diminishes the value of each unit of foreign currency in terms of the local one. First of all

[8] See Krusell et al. (2000) for a discussion over this restriction and its plausibility.

the shock reduces the revenues of the exporting firms by $ys_x\frac{\Delta e}{e}$, where s_x denotes the share of exported output.[9] If firms are at the top of the quality ladder (Grossman and Helpman, 1991a), they might reduce the price and still keep on producing and selling the same amount of output. Alternatively, they might find convenient reducing the production, rescaling the entire production chain. If all inputs are flexible this scale effect would impact labor demand of highly skilled and lower skilled workers symmetrically.

The reduction in the relative price of offshored labor increases the import of semiprocessed materials and has a negative effect on low-skill labor. In addition, it lowers the total price of the composite input X with positive effects on high skill labor and imported capital $(1/(1-\sigma) < 1)$. The shock impacts also the price that Swiss firms pay to import foreign capital. This generates an increase in increase in the demand for high skill labor because of the complementarity amongst the two. It also impacts positively the demand for low-skills labor and offshored labor, but the increment is relatively lower than the one on high-skill labor because of the restriction $1/(1-\rho) < 1/(1-\sigma)$.

In this framework, a reduction in import prices impacts the labor demand through both channels in the same direction: it generates a shift from workers performing tasks that can be substitute by foreign capital or offshored to highly skilled workers complementing the new imported machines. Motivated by the implications of this simple model, in section 1.5 I develop a pre-shock measure of exposure to substitutability accounting for both channels and study how firms with different exposure reacts to the price shock.[10]

1.3 Data

1.3.1 Data Sources

I rely on three sources of data. The first are transaction-level records of Swiss imports and exports collected by the Swiss Customs Administration (Swiss-Impex). This data

[9] The average output selling price is $(1-s_x)p^D + s_x e p^X$.

[10] Previous literature shows routine tasks are easier to be substitute by machines (Autor et al., 2003; Acemoglu and Autor, 2011), therefore I use the routine intensity indicator of the labor force to account for exposure to imported capital. I serve by the offshorability index constructed by Autor and Dorn (2013) to account for the offshoring channel.

starts in 2014 and covers all goods that enter Switzerland or leave the country on a daily bases. For each good it reports the value-at-the-border in Swiss francs and the Harmonized-System 8-digit product classification. In addition, for each transaction the data records detailed information on the recipient and the sender, including the name and the exact address of the origin and destination and unique identifiers for importing and exporting firms (UID).[11]

The firm identifier allows me to merge the import-export information to the second source of data: the near-universe of online job postings in Switzerland scraped and assembled by a private labor market analytics firm. This data contains information such as name, industry and market of the firm or the recruitment agency that posted the job ad, the job location, the posting and closing dates, the position, the quota (full time/part time), the language of the advertisement, the job title and the full description of the job. The database covers the period from 2012 to 2021 for an average of almost 1 million job ads per year issued by 105,848 firms or public entities and 3,122 headhunters. The geographical representation closely follows the employment distribution of Switzerland with 19% of all postings associated with the Canton Zurich and 10% with the Canton Bern, followed by Argovia, St Gallen, Lucern and Vaud. Additional information on the job requirements are extracted from the title and the description by means of textual analysis. This additional information includes: occupation (job), skills, education, experience and language. The skill content is the main variable of interest which allows for a detailed analysis of skill demand within each occupation: a margin that has been often ignored in the literature, but hass recently been shown to be relevant by Hershbein and Kahn (2018) and Deming and Kahn (2017). More than 10,000 skills were classified. Each job ad requires, on average, 3.4 skills. In addition, a second skill classification is extracted: soft skills, accounting for 144 distinct categories and 4.5 entries per job ad, on average.

The last source of data is Bureau van Dijk's ORBIS database, which contains publicly available information about firms, including balance sheets and income statements, as well as a detailed description of the activity of the firm.

[11] The firm identifier is not included for the years before 2016. I generate it through the firm name and a cross-walk name-UID constructed using the subsequent periods.

1.3.2 Descriptive Statistics

Following the previous literature on trade and the labor market, I focus only on manufacturing firms. This is also the sector with the strongest trade ties to the rest of the world, while it employs a substantial part of the Swiss labor force. In order to focus on firms active in the labor market, I restrict the sample to firms posting at least 5 job ads in the period 2012-2014 and at least 10 job ads in the period 2012-2017. The result is a list of 1,752 firms of which 89% made at least one import transaction and 81% at least one export transaction in the period 2014-2015.

I use the country of origin/destination to separate transactions with the Euro area from the rest fo the world. I merge product classifications with the Broad Economic Categories (BEC) to distinguish consumption goods from intermediate goods and capital goods. I translate all the skills to English and merge them with the Burning Glass Technology skill classification system, which is largely used throughout the literature. This skill-matching allows me to group skills into 512 clusters and into, more tractable, 28 skill families.[12] Using the job titles, and following the algorithm created by Mihaylov and Tijdens (2019), I generate five task measures related to each job ad: non-routine analytic, non-routine interactive, routine cognitive, routine manual and non-routine manual tasks. The five measures sum up to 1 for each job and are then aggregated into a routine intensity index (RTI) ranging between -1 and 1: 1 indicates that the occupation contains only routine tasks, and -1 indicates that the occupation contains only non-routine tasks.[13] Still using the job titles, I also assign to each vacancy an offshorability measure based on the index constructed by Autor and Dorn (2013). This indicator assigns to each occupation a value based on the degree to which the job can be relocated abroad. In particular, it indicates *how much* a job requires interpersonal interac-

[12]Figure 1.C1 in appendix shows two examples of job postings together with the skills and the skill families extracted. Figure 1.D2 in appendix exhibits the distribution of the variables *skill* and *skill families*. Figures 1.E1 and 1.E2 in the appendix highlight the presence of substantial variation in the distributions of skills by industry and occupation, respectively. This highlights how an analysis at the skill level can point out dynamics that can not be analyzed using traditional measures of occupations or industries.

[13]This classification follows the idea of conceptualizing jobs as a series of occupational tasks which is labeled as the "task-based approach" (Autor et al., 2003; Autor, 2013; Autor and Dorn, 2013). Each occupation is divided into different tasks, for a total of 3,264 occupation-specific tasks, and each task is classified according to each of the five categories. These values are then averaged over all the tasks that are performed in each occupation. The five measures were shown to be positively correlated with the equivalent measures used in Acemoglu and Autor (2011) at the occupational level, with correlation coefficients ranging between 0.41 and 0.70. The RTI is just the difference between the routine and the non routine indexes.

tion or high proximity between the worker and the work location to be performed.[14]

I collapse all the information at firm-month level. If a firm does not appear in the customs data for a given month, I impute zero imports and export. Similarly, if a firm does not show up in the job posting data for an entire month, I infer that the firm did not post any job advertisement, while the skill requirements measures are registered as missing for that given month. The constructed database contains, among others: information on total import and export by type of good and destination country; the number of new job advertisements by skill; the average RTI; minimum education requirement; characteristics of the firm such as industry, location, yearly information on capital; operating revenue and number of employees.

In table 1.1, I report descriptive statistics for the main variables for the period 2014-2016. Panel A shows characteristics of the vacancies posted. The average monthly number of new job ads posted is 1.27 per firm, for a total of 79,860 postings. Only 8.7% of the job ads posted requires a bachelor degree.[15] The average routine index in the sample is -0.487, showing a tendency towards labor demand focused mainly on non-routine jobs. A standard deviation of 0.62 together with the boundaries values show that there exists a significant variation in this indicator, that can be analyzed. The job offshorability measure ranges from -3.011 to 2.661 and has an average of 0.391 and a very large variance with a standard deviation of 1.025.

Panel B reports the amount of monthly import and export transactions made by the firms in the sample. Total annual exports for firms in the sample account for 38.6% of the total Swiss exports in 2014. The sample also captures roughly 22.62% of the total imports in the same year. This is not surprising since exports are more concentrated amongst bigger firms and imports also account for a vast portion of private consumption not related to firms. Slightly more than 50% of the trade in the sample is between Switzerland and the Euro Area - which is directly affected by the policy shock. Finally, I divide import and export into four groups, according to the currency with which transactions were made[16]. For both imports and exports most of trading is invoiced in euros. While one-third of exports concerns transactions in Swiss francs, less than 20% of

[14]The indicator is based on two variables derived by Firpo et al. (2011) using the US O*NET database: *Face-to-face contact* and *on-site job*.

[15]For the remaining portion, it is not possible to state the exact required education.

[16]Bonadio et al. (2019) highlight important differences in the pass-through effect of the Swiss franc shock on prices dependent on whether the transaction was invoiced: for goods invoiced in euros, the pass-through is immediate and complete while for goods invoiced in Swiss francs it is slowed and partial.

Table 1.1: Summary statistics

Variable	**Mean** (1)	**Std. Dev.** (2)	**Min.** (3)	**Max.** (4)	**N** (5)
Panel A: Job Postings - by job ads					
Bachelor Degree or Higher	0.087	0.282	0	1	79860
RTI	-0.487	0.62	-1	1	72943
Job Offshorability	0.391	1.025	-3.011	2.661	72055
Panel B: Transactions - monthly					
Total					
Total Monthly Import	3.133	50.026	0	2183	63072
Total Monthly Export	5.569	77.523	0	3054	63072
By Type of Goods					
Monthly Import Intermediate	2.534	48.02	0	2183	63072
Monthly Import Consumption	0.129	1.42	0	192	63072
Monthly Import Capital	0.153	1.086	0	108	63072
Monthly Export Intermediate	3.531	63.39	0	2767	63072
Monthly Export Consumption	0.297	4.305	0	223	63072
Monthly Export Capital	0.581	3.188	0	153	63072
Euro Area					
Monthly Import Euro Area	1.586	28.246	0	1771	63072
Monthly Export Euro Area	2.507	36.38	0	1717	63072
By Currency					
Monthly Import CHF	0.492	17.783	0	1454	63072
Monthly Import EUR	1.417	17.6	0	883	63072
Monthly Import USD	1.082	29.353	0	2017	63072
Monthly Export CHF	1.773	29.716	0	2198	63072
Monthly Export EUR	1.989	30.287	0	1708	63072
Monthly Export USD	1.03	21.009	0	1224	63072
Panel C: Firm Information - yearly					
N. of Employees	209.964	3108.477	1	133413	5100
Capital (in million CHF) in	4.598	41.573	0.02	1228	5068
Operating revenue (in million CHF) in	157.65	2020.389	0.25	60983.953	2309
Listed	0.01	0.098	0	1	5256

Notes: This table shows the monthly summary statistics of some job ad variables in panel A, trade variables (in million CHF) on a monthly basis in Panel B and firm information on a yearly basis in Panel C for the period 2014-2016. The number of observations in panel B (63,072) is the product of 1,752 firms × 12 months × 3 years. The firm information (Panel C) is available only for a portion of the sample: number of employees and capital are present for about 97% of the firms, operating revenues for less than half of the sample. Number of employees and operating revenues are imputed for about 98% of the sample. Total Swiss Export 2014: 285,079 million CHF - Total Swiss Import 2014: 282,505 million CHF. Trade made by firms in the sample in 2014 accounts for 38.6% of total Swiss exports and 22.62% of total Swiss imports in 2014. Total Swiss employment in 2014: 4,897 million workers. In 2014, the firms in the sample employ 7.5% of the total Swiss employment.

imports are invoiced in Swiss francs. Conversely, the US dollar is more prominent for importing transactions.

Panel C shows basic firm characteristics extracted from ORBIS for a subsample of firms for which data is available. These firms have an average operating revenue of 157 millions Swiss francs and account for 7.5% of the Swiss labor force.

1.4 Trade and Skill Requirements

In this section, I show how pre-shock values of trade and skill requirements are related and how they change after the shock. In section 1.5, I report results of a more complex analysis aimed at identifying how firms respond to the shock in terms of trade and how these adjustments impact their labor demand.

1.4.1 Cross-Sectional Correlation prior to 2015

I start by considering only the year 2014; a year of relative stability in terms of the global and the Swiss economies, and with an (artificially) stable exchange rate between the Swiss franc and the Euro. I divide the firms in my sample into four groups according to their level of trade. In particular, I consider two measures: total imports per worker and total exports per worker. Since the resulting distribution of firms along these measures are highly skewed, I sort firms over these measures and group them into four clusters based on quantities of the respective distributions.

Table 1.2 reports average skill requirements by imports per worker in panel A, and exports per worker in panel B. By construction each group contains an equal number of firms. The figure in table 1.2 shows that labor demand is highly correlated with trade. Consistent with the theory that trading firms employ less workers in routine jobs, I find that the routine intensity indicator monotonically decreases as the imports per worker increases. The gap between firms in first and fourth quartiles is 0.12 points, equivalent to a reduction of about 19% in the routine task intensity. A similar pattern arises for exporting firms: the change between the top and bottom quartiles is even greater in this case, however, there is less variation in the middle of the distribution. This result can be due to the high correlation between the two trading measures as well as to other channels, aside from offshoring, through which trade impacts on skill requirements.

Table 1.2: Summary statistics - Trade

						Weighted		
	N. Firms	**N. ads**	**RTI**	**Manuf.**	**I.T.**	**RTI**	**Manuf.**	**I.T.**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Import per worker								
First quartile	423	.837	-.306	.152	.378	-.374	.124	.397
Second quartile	423	.825	-.309	.127	.405	-.377	.109	.386
Third quartile	423	1.593	-.354	.093	.421	-.476	.086	.487
Fourth quartile	424	2.406	-.429	.097	.466	-.501	.083	.487
B. Export per worker								
First quartile	423	.836	-.283	.125	.363	-.342	.107	.388
Second quartile	423	.997	-.329	.109	.393	-.372	.088	.392
Third quartile	423	1.342	-.334	.112	.42	-.468	.096	.443
Fourth quartile	424	2.485	-.443	.109	.481	-.524	.089	.518

Notes: This table reports the average number of job postings per month, the average routine Intensity Indicator and the share of ads requesting Manufacturing and I.T. skills for four groups of firms in 2014. Firms are grouped based on total yearly import per worker in Panel A, and export per worker in Panel B. Averages in columns (6-8) are weighted by the number of ads posted.

To further investigate the relation between trade and selected skill requirements, columns (4) and (5) of the table look at two families of skills: manufacturing skills and information technology (IT) skills, respectively. The choice of these two families of skills is driven by the nature of the labor demand I analyze. The sample is composed by manufacturing firms, which traditionally employ workers with manufacturing skills. If outsourcing or automation take place, the firm might rather employ individuals with skills that are not closely related to the main firm activity. IT skills here are intended to be understood in a broad sense, this class includes all the job advertisements that mention a specific software or a programming language, ranging from *data management* to standard software use skills like *microsoft office* skills, and more complex programming language skills, for example, *C++* and *Java*.[17] A clear pattern arises: firms in the fourth import quartiles require 36% less manufacturing skills and 23% more IT skills with respect to firms in the first quartiles. While for IT skills the gap is slightly greater in the export case, for manufacturing it is about half. This difference might be explained by the nature of these requirements and the mechanisms behind the two forces. Both international buying and selling require some kind of communication infrastructure which can be performed by workers with IT skills. On the other hand, offshoring and the input of production factors from abroad substitutes mainly for workers with manufacturing skills. Columns (6), (7) and (8) show that the patterns highlighted above also remain stable when using weighted measures of skill requirements as outcomes. Table 1.F1 in the appendix shows linear correlations between trade and skill indicators, confirming these patterns.

1.4.2 Evolution of Trade and Skill Requirements

Panel A of figure 1.2 reports the total quarterly exports and imports for Swiss firms during the period 2014-2016. It shows an initial stable evolution of exports and a slight decrease in imports in 2014 and in the first two quarters of 2015. Starting from the last quarter of 2015 there is an increasing and roughly parallel trend in both measures.

There is high volatility in the data, and the distribution of imports and exports is

[17] Tables 1.B4 and 1.B5 in the appendix provide a description of the skill clusters associated with these skill families. I also regress imports and exports on all the skill families in the sample, highlighting positive correlation between imports and IT skills and a negative one between imports and manufacturing is negative.

Figure 1.2: Evolution of Trade

(a) Log of Total Imports and Exports (b) Average Log Imports and Log Exports

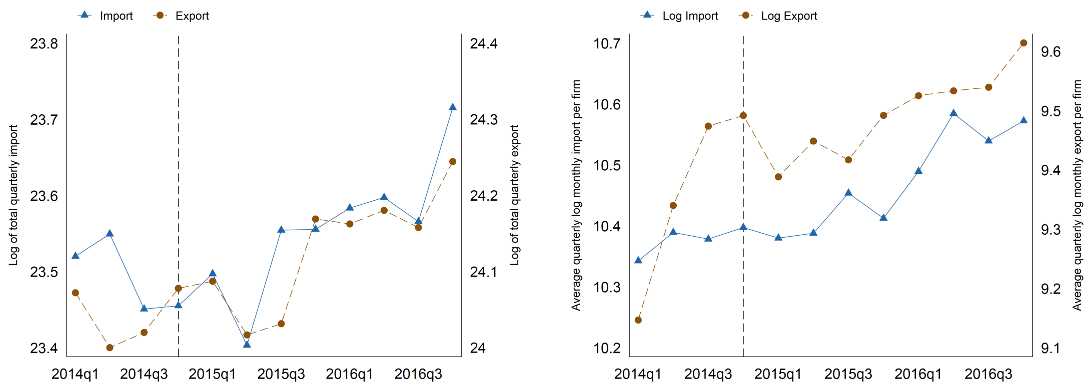

Notes: This figure shows the amount of imports and exports for Swiss firms from 2014 to 2016. Panel A reports the logarithm of the total amount of imports and exports generated in each quarter by all the firms. Panel B reports the average of the monthly log imports and exports generated by each firm.

strongly skewed toward the right.[18] For this reason, I also consider log monthly imports and exports and I plot their quarterly averages in panel B of figure 1.2. This measure is less sensitive to outliers and better describes the behavior of firms in the middle of the distribution. In contrast to the trends in panel A, the quarterly averages plotted in panel B imports and exports evolve much more smoothly. It is indeed possible to identify a stable trend in imports in 2014 and most of 2015, followed by a rise towards the end of the time period plotted. The opposite is true for exports: the amount of exported goods from Switzerland constantly increased in 2014. This trend ended after the shock and total exports remained relatively constant in 2015 and 2016. The trend inversion for both variables is plausibly driven by the exchange rate shock, which generated a change in prices making imports less costly and exports more expensive.

In figure 1.3, I report the evolution of labor demand (panel A) and several measures of skill requirements (panels B-F). The number of new jobs posted online by the firm in the sample, constantly decreases in 2014. This unusual declining pattern is consistent with employment in the Swiss context.[19] Labor demand remains stable in 2015 and increases slightly in 2016. The number of total skills identified in the text of job advertisements constantly increases (panel B). It goes from approximately 3.1 skills re-

[18]Figure 1.D1 in the appendix shows the distribution of monthly imports and exports highlighting the skewness of the distribution and the concentration of most of the exports to fewer firms, rather than imports, which is less concentrated.
[19]See www.bfs.admin.ch for a description of Swiss employment in the last decade.

Figure 1.3: Evolution of Labor Demand in time

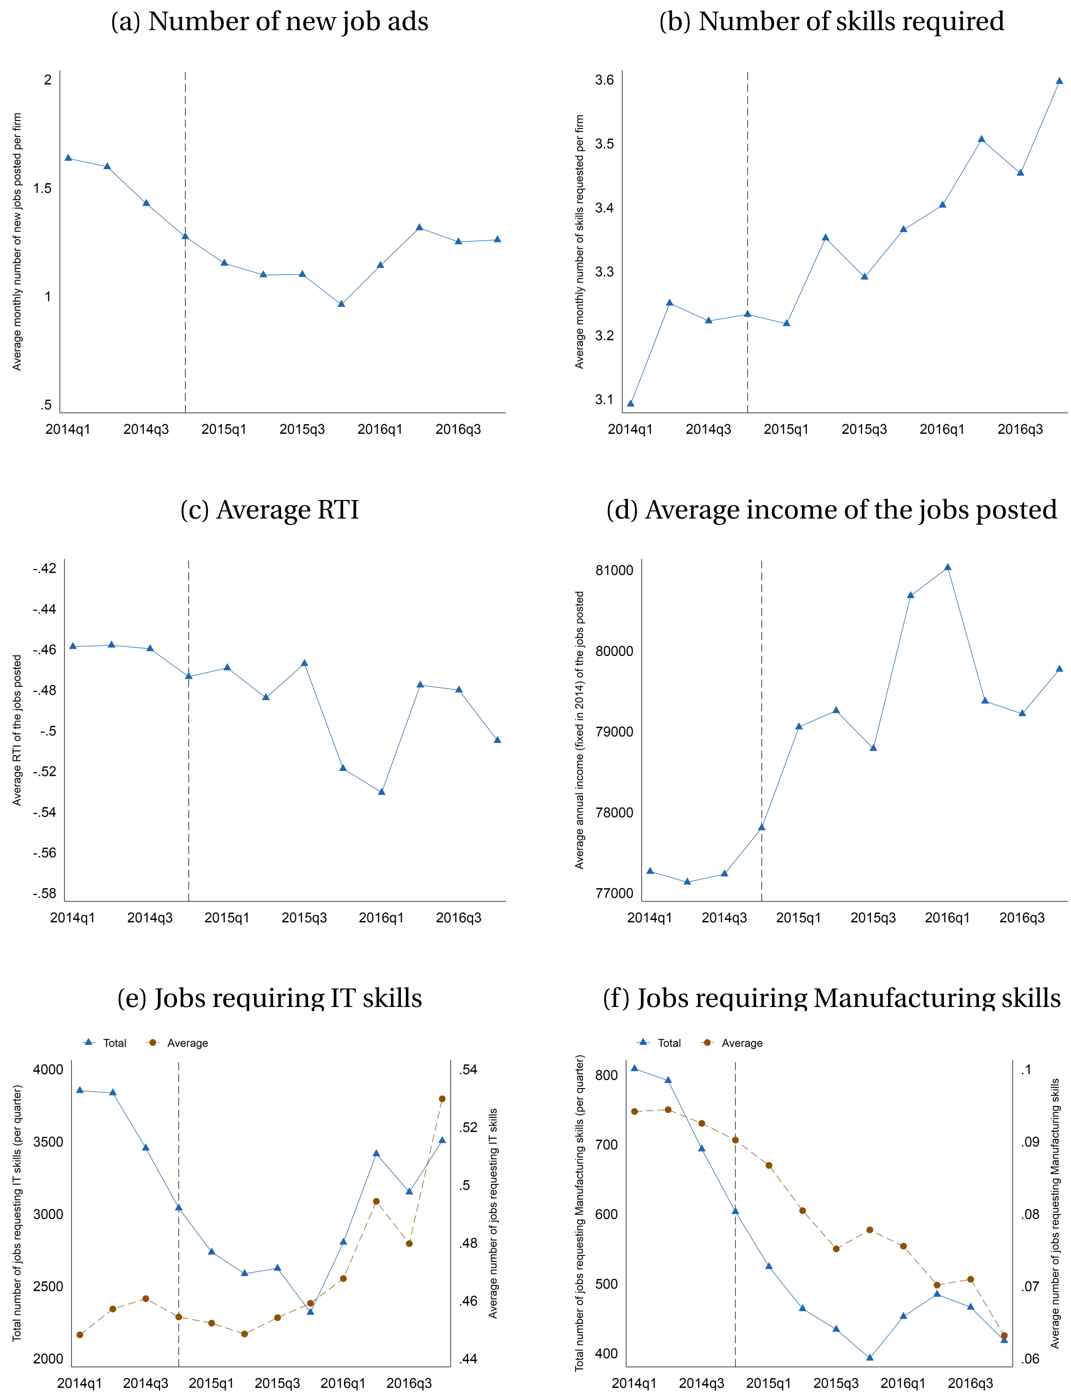

Notes: This figure shows the average monthly new job postings (panel A) the average number of skills requested (panel B), the average RTI (panel C) and the average income of the posted jobs (panel C) from 2014 to 2016, by quarter. Income is calculated at the occupational level in 2014. Panel E and F plots the number, and the average, of job advertisements requiring IT skills and manufacturing skills for the same period, by quarter.

quired per job in the first quarter of 2014 to 3.6 in the last quarter of 2016. Panels C and D show the evolution of the routine intensity index (RTI) and the annual income associated with each job posted.[20] To avoid mixing composition effects with change in salary, I consider the income of each occupation as fixed at its level in 2014. Therefore, any change has to be addressed only to a variation in the composition of labor demand in terms of jobs. The RTI does not vary much. It is quite stable around -0.46, with a slight decrease at the end of 2015 and the beginning of 2016, as well as in the last quarter of 2016. The income measure shows a constant increasing pattern starting in the first quarter of 2015 and ending in the first quarter of 2016. The range is about 3,000 Swiss francs per year.

Panels E and F of figure 1.3 plot the total and the average number of job advertisements requiring IT and manufacturing skills. As a consequence of the reduction in the number of job postings in 2014, both measures show a declining pattern for the first six quarters. However, when moving to the average values two quite mirroring trends arise. The trend in IT skills is constant in the first half of the period and almost monotonically increases starting in the third quarter of 2015. Conversely, manufacturing skills show a mild decreasing trend in 2014 and decline more rapidly in 2015 and 2016. These changes in skill requirements, taken together with the increasing imports reported in figure 1.2 signal a potential transformation of the firms' production functions following the exchange rate shock, substituting imported inputs and manufacturing skills. I investigate this causal link further in the next section.

1.4.3 Post shock Changes in Trade and Skill Requirements

Earlier in this section, I show (i) the existence of a strong correlation between trade and different measures of skill requirements and (ii) an increase in monthly imports, in the demand for jobs with higher payment and in the requirements for IT skills, as well as a decline in the postings targeting manufacturing skills after the shock. In order to establish a link between changes in trade after the shock and changes in skill requirements, I first group firms based on the change in log imports and log exports per worker between 2014 and 2015. I then show how labor demand in 2015 is distributed

[20] Annual income per occupation is extracted from workers surveyed by Swiss labor force survey in 2014.

among these groups of firms.

Results are reported in table 1.3. Firms in the second and third quartiles in both panels A and B posted a greater number of job advertisements in 2015. Panel A shows that the greater the increase in imports after the exchange rate shock, the lower the average routine intensity indicator of the jobs that are posted. This is valid for the first three quartiles, but not for the last quartile. Changes in imports do not seem to have an impact on labor demand instead. Manufacturing skills are strongly negative correlated with the change in imports, while IT skills positively correlate with this change. The first-fourth quartile change for manufacturing skills is sizeable: a 53% reduction. The IT skills change is also important, but less significant, as it accounts for a 12% increase. Additionally, as showed in panel B, the change in exports does not seem to be strongly correlated with these two measures of skill requirements. This suggests that imports, and their substitutability with low skilled workers, is the channel through which the exchange rate shock affected the composition of labor demand.

Table 1.3: Changes in Trade and Skill Requirements

						Weighted		
	N. Firms	**N. ads**	**RTI**	**Manuf.**	**I.T.**	**RTI**	**Manuf.**	**I.T.**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Δ Import per worker								
First quartile	391	.646	-.348	.128	.411	-.404	.109	.423
Second quartile	391	1.076	-.382	.099	.423	-.474	.089	.491
Third quartile	391	1.734	-.41	.082	.456	-.534	.068	.488
Fourth quartile	392	.852	-.373	.081	.459	-.435	.077	.469
B. Δ Export per worker								
First quartile	354	.877	-.385	.087	.464	-.482	.079	.507
Second quartile	354	1.288	-.388	.105	.463	-.492	.089	.514
Third quartile	354	1.603	-.401	.103	.454	-.484	.076	.451
Fourth quartile	354	.798	-.387	.085	.412	-.377	.077	.428

Notes: This table reports the average number of job postings per month, the average routine Intensity Indicator and the share of ads requesting Manufacturing and I.T. skills for four groups of firms in 2015. Firms are grouped based on the difference between the log total yearly imports per worker in 2015 and the log total yearly imports per worker in 2015 in Panel A, and on the difference between the log total yearly exports per worker in 2015 and the log total yearly exports per worker in 2015 in Panel B. Averages in columns (6-8) are weighted by the number of ads posted.

1.5 Empirical Analysis

Results in the previous section show that firms which increased their import shares changed their demand for skills in a way consistent with the theory on the substitutability between low skills and imported goods. To further establish a link between exchange-rate-driven changes in imports and changes in skill requirements, in this section I classify firms based on their exposure to the substitutability of their labor force and then estimate a series of ordinary least squares (OLS) and two stage least squares (2SLS) models to understand how firms with different levels of exposure reacted to the shock.

1.5.1 Exposure Measure

The price effect generated by the exchange rate shock potentially affected the entire Swiss economy. However, not all firms could react to it and not all by the same magnitude. The reaction depends on which portion of the production chain a firm can offshore and which share of workers it can substitute with imported inputs. To differentiate among different types of firms, I construct a measure of *exposure to substitutability* by firm. This new measure combines the offshorability index and the routine intensity indicator of pre-shock firm labor demand, and aims at indicating how much the labor force of a firm can be substituted either by the relocation of some activities abroad, or by the import of capital from abroad.

Swiss manufacturing activity is unique (section 1.2.1), as it is highly concentrated on high standard products, located at the top of the quality ladder (Grossman and Helpman, 1989, 1991a,b). In some sectors, Swiss producers are quasi-oligopolist in the world market and use their privileged position to set a high quality-adjusted market price for their products. In most cases, the *Swiss Made* label associated with their products is essential to guarantee high mark-ups when placing the goods on the market. Offshoring can therefore be less attractive to Swiss firms. For example, jobs like *handicraft workers in textile, leather and related materials, metal moulders and core-makers* and *weighers* reside in the top decile of the offshorability index, but it is very unlikely that firms employing these workers relinquish their privileged position in the market by offshoring their activities. However, instead of offshoring, firms can replace part of their labor force through automation, mantaining the production in Switzerland, while benefiting from the cheaper price of imported capital from abroad generated by the exchange rate shock.

To capture both channels through which the currency shock might impact labor demand in the Swiss context, I augment the offshorability index with the RTI. The new *substitutability* measure consists of the average between the offshorability index and the routine intensity indicator.[21] By including these two indexes, this new substitutability measure is intended to indicate the potential degree to which a job can be substituted either by a machine or by foreign workers.

I start by assigning an offshorability index value Off_{jit} to each vacancy j posted by firm i at time t through the job title.[22] I then create rescaled measures ranging between 0 and 1 of the offshorability ($StOff_{jit}$) and the routine intensity index ($StRTI_{jit}$) of the vacancy j, and average them with equal weights to construct the *substitutability* index (sub_{jit}), as shown in equation 1.3.[23]

$$sub_{jit} = \left(StOff_{jit} \ + \ StRTI_{jit}\right)/2 \tag{1.3}$$

I then use the job advertisements posted in the 2 years before the analyzed period, 2012 and 2013, to construct a measure of exposure to substitutability for each firm (sub_i). This is simply the average of the substitutability of all the job advertisements posted by firm i in those two years.

$$sub_i = \sum_j \frac{sub_{jit}}{N_{it}} \qquad \forall\, t \in [2012, 2013] \tag{1.4}$$

Table 1.4 reports descriptive statistics in 2014 for the exposure measure, the two variables used to construct it, and the main variables used in the analysis. The average level of exposure is 0.44 with a standard deviation of 0.12. Panel (B) and (C) of the same table reports descriptive statistics for exposed firms, that is, firms with an expo-

[21]The relation between routine and offshorability index is not straightforward: in a survey, Blinder and Krueger (2013) find that the two measures are non correlated, while Autor and Dorn (2013) find a correlation index between routine and offshorability across US commuting zones of 0.66. In this database the correlation between the two indicators at the occupational level is also positive (ρ=0.19).

[22]The index was constructed following the SOC occupation nomenclature. In order to assign a value to each job in the database I transformed and adapted it to the ISCO-08 occupation classification using the crosswalk document provided by the US Bureau of Labor Statistics.

[23]Figure 1.D3 in the appendix presents the distribution of the three variables.

Table 1.4: Summary statistics - Variables in the main analysis

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**	**N**
	(1)	(2)	(3)	(4)	(5)
Panel A: *All Firms*					
Exposure to Substitutability	0.443	0.116	0.114	0.961	1653
RTI Index	-0.309	0.46	-1	1	1653
Offshorability Index	0.324	0.654	-2.629	2.661	1655
Manufacturing Skills	0.125	0.228	0	1	1627
IT Skills	0.394	0.33	0	1	1627
Log Import	10.92	4.086	0	21.001	1752
Log Export	10.252	5.274	0	21.564	1752
Panel B: *Firms with exposure lower than the median*					
Exposure to Substitutability	0.356	0.058	0.114	0.428	827
RTI Index	-0.43	0.418	-1	1	776
Offshorability Index	0.243	0.635	-2.629	1.963	778
Manufacturing Skills	0.117	0.223	0	1	761
IT Skills	0.416	0.331	0	1	761
Log Import	11.249	4.188	0	21.001	827
Log Export	10.759	5.315	0	21.564	827
Panel C: *Firms with exposure greater than the median*					
Exposure to Substitutability	0.53	0.092	0.429	0.961	826
RTI Index	-0.197	0.467	-1	1	780
Offshorability Index	0.407	0.652	-1.934	2.661	781
Manufacturing Skills	0.129	0.228	0	1	771
IT Skills	0.382	0.327	0	1	771
Log Import	10.643	4.035	0	20.038	826
Log Export	9.870	5.24	0	19.864	826

Notes: This table reports summary statistics for vacancies posted in 2014. Each observation is a firm. Exposure to Substitutability is constructed using RTI and Offshorability in the period 2012-2013. Manufacturing and IT Skills are skill (families) requirements extracted from the text of job advertisements. Log Import and Log Export refer to the logarithm of the monthly average over each year.

sure measure greater than the median, and the remaining unexposed firms. By construction, the substitutability measure as well as the RTI and the offshorability index are substantially lower for unexposed firms. Additionally, the other skill requirements indicators and trade variables present some pre-shock differences consistent with the predictions: exposed firms request more manufacturing and less IT skills, and present lower values of trade.

Figure 1.4 presents the distribution of the exposure measure across firms by subsector. The substantial variation within each industry displayed in figure 1.4 highlights the presence of firm-specific behaviors that can be detected and investigated with the firm-level data used in this paper.

Figure 1.4: Distribution of the exposure by sub sector

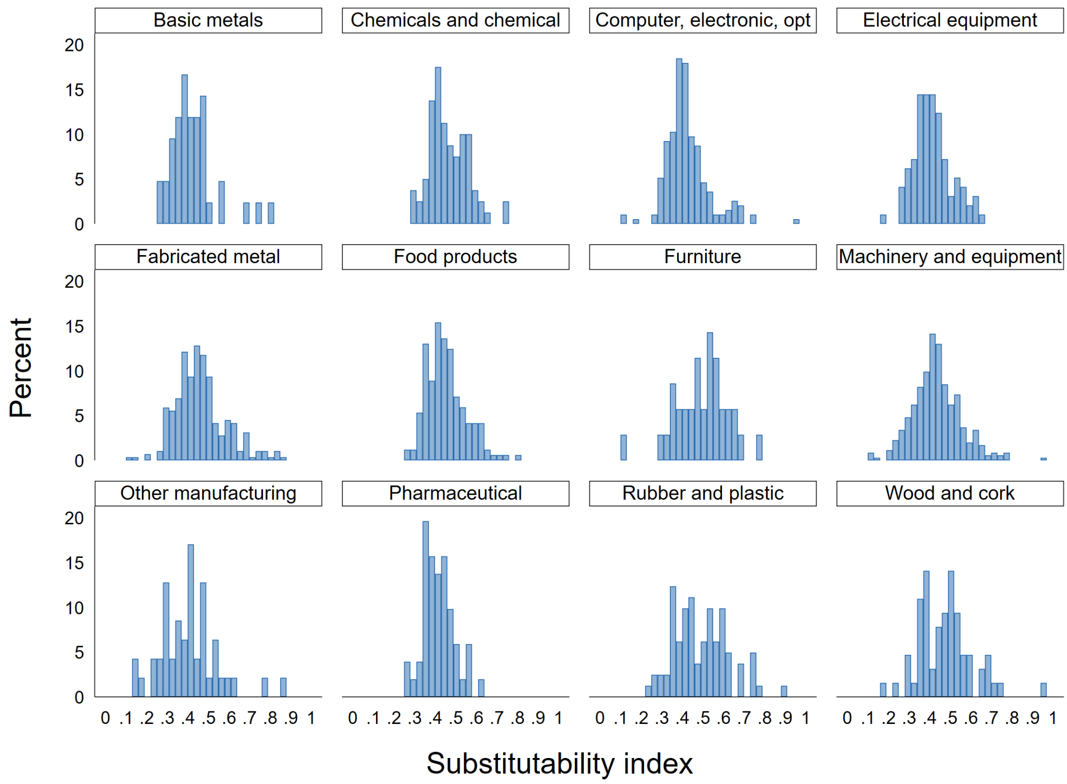

Notes: This figure shows the the distribution of the exposure index by sub sector for the period 2014-2015. Each observation is a firm-year cell. Only sectors with more than 100 observations.

1.5.2 Measuring Effects on Skill Requirements

I begin by using a simple difference in differences framework to measure how firms with different levels of exposure to substitutability change their labor demand after the shock. I estimate the following model.

$$Y_{it} = \alpha_1 + \gamma_{1,i} + \lambda_{1,t} + \delta_1 X_{it} + \beta_1 sub_i \times D_{t=post} + \varepsilon_{it} \tag{1.5}$$

Y_{it} is a measure of skill requirement for firm i at time t, $\gamma_{1,i}$ and $\lambda_{1,t}$ are firm and time fixed effects, respectively, X_{it} is a vector of firm time variant characteristics, $D_{t=post}$ is an indicator taking value 1 for the post ban period and sub_i is the exposure to substitutability described above. The coefficient of interest is β_1, which is intended to estimate the effect of the shock on skill requirements for a unitary change in the exposure measure.

I collapse the data at the firm-year level and fit the model in equation 1.5 for three outcomes. Table 1.5 presents the results.

Table 1.5: Skill Requirements - Results

	OLS Estimation					
Dependent Variable:	*RTI*		*Manuf. Skills*		*IT Skills*	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.035	0.067	0.015	0.032	-0.019	-0.011
	(0.058)	(0.041)	(0.029)	(0.022)	(0.044)	(0.033)
Exposure (sub)	1.059***		0.073		-0.052	
	(0.101)		(0.051)		(0.077)	
Exposure × Post	-0.175	-0.196**	-0.056	-0.086*	0.090	0.069
	(0.127)	(0.089)	(0.064)	(0.048)	(0.097)	(0.073)
Controls	✓	✓	✓	✓	✓	✓
Firm FEs		✓		✓		✓
N	4232	4232	4121	4121	4121	4121

Notes: This table shows estimated coefficients of equation 1.5. Dependent variables are: the RTi index in columns (1) and (2), the share of Manufacturing Skills in columns (3) and (4), and the share of IT skills in columns (5) and (6). All variables are year averages. Capital and Number of Employees are included as controls. *Post* is a dummy taking value 1 in years 2015 and 2016.
Robust standard errors in parentheses. Table 1.G2 in the appendix presents estimates with standard errors clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Columns (1) and (2) report the results on the routine intensity indicator (RTI), which ranges between -1 and 1. The exposure measure strongly correlates with RTI in the pre-shock period: the estimated coefficient of 1.059 indicates that for a unitary increase of the exposure measure, the RTI increases by 1.059. Considering that the exposure measure is bounded between 0 and 1 and is constructed, for its 50%, using a scaled RTI measure, this coefficient shows substantial consistency between job ads posted in 2012-2013 and those posted in 2014. The interaction coefficient is negative, signaling a decrease in the RTI of jobs posted after the shock for firms with high level of exposure. The currency shock lowers the role of exposure, by 0.175 and once firm effects are included, by 0.195. The shock weakens the link between exposure and RTI: a unitary increase in the exposure lowers the routine intensity of the labor demand by 14.7%.[24]

The lower the effect of the shock on the demand for manufacturing skills (column 4), the greater the exposure to substitutability. In particular, an increase of 1 in the exposure measure translates into a decrease in the share of the job postings requesting manufacturing skills by 0.086, which equals to a 74% increase. In contrast, I do not find any difference in the share of job ads requesting IT skills by exposure (column 6). However, the sign of the interaction coefficient is positive signaling a positive correlation between the exposure and the change in IT skills after the shock.[25]

To provide further evidence that the effect is the result of the currency shock and is not due to any confounding trend, I estimate the event-study version of the results in table 1.5. I separately regress each indicator of skill requirements Y_{it} on firm (μ_i) and quarter (ζ_q) fixed effects, firm specific controls (X_{iq}) and interaction terms between quarter dummies ($D_{q=l}$) and the exposure measure (sub_i).

$$Y_{iq} = \kappa + \mu_i + \zeta_q + \eta X_{iq} + \sum_{l=2014,q1}^{2016,q4} \rho_l \, sub_i \times D_{q=l} + \xi_{iq} \tag{1.6}$$

Figure 1.5 shows that there are no significant treatment effects of the exposure on skill requirements in the pre-shock period and provide information on the *speed* of the effect. The effect on RTI constantly increases for the first three quarters and then

[24] This result is computed taking into account the distribution of the main variables of interest reported in the appendix table 1.4

[25] Table 1.G2 in appendix shows that these results are robust to clustering standard errors at firm level. Figure 1.H1 in appendix shows that exposed firms post also more job advertisements just after the shock. This effect vanishes 6 quarters after the shock signaling a sort of adjustment period in which the firm adapted its workforce.

stabilizes, while for maunfacturing skills, the effect becomes stables after two quarters.

Figure 1.5: Event Study - Skill Requirements

(a) ES - RTI

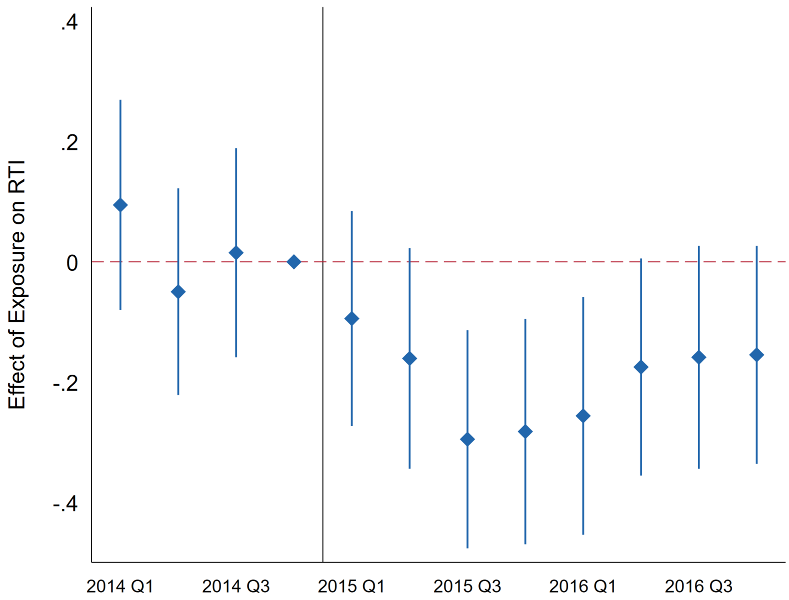

(b) ES - Manufacturing Skills

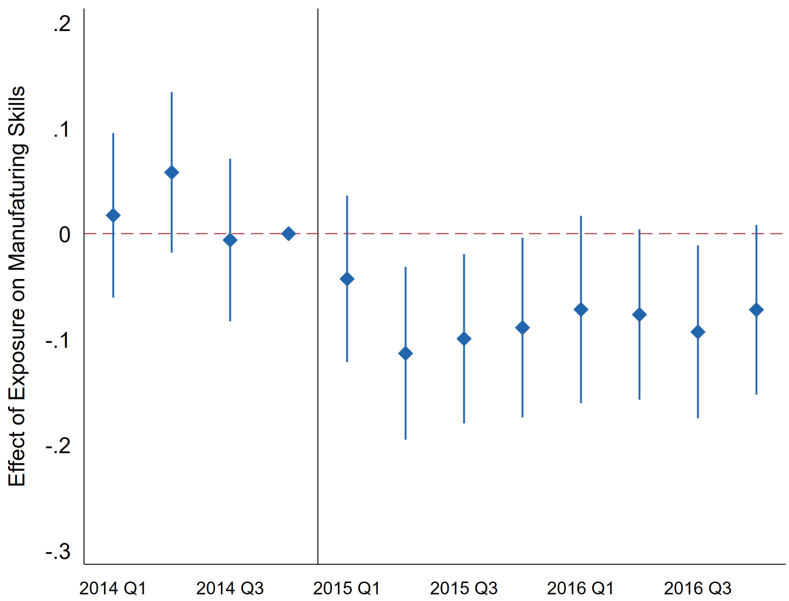

Notes: The figure plots the estimated treatment effects ρ_l in equation 1.6, and their confidence intervals, for RTI in panel (a) and Manufacturing skills in panel (b).

I further perform a placebo exercise. I compute the exposure measures using only information in 2012, and re-estimate the model in equation 1.5 over the period 2013-2014, modeling a placebo shock in 2014.[26] Results in table 1.6 show no changes in skill requirements after the shock for exposed firms. The interaction coefficients for RTI and manufacturing skills are not significant and in terms of magnitude are four times smaller than those identified in table 1.5; the placebo treatment effect on IT skills is also not significant and its sign is opposite to that found in table 1.5.

Table 1.6: Skill Requirements - Placebo

OLS Estimation						
Dependent Variable:	*RTI*		*Manuf. Skills*		*IT Skills*	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.051	0.048	0.030	0.013	0.025	0.014
	(0.061)	(0.053)	(0.035)	(0.024)	(0.052)	(0.039)
Exposure (sub)	0.878***		0.104*		-0.079	
	(0.095)		(0.056)		(0.082)	
Exposure × Post	-0.054	-0.045	-0.056	-0.022	-0.044	-0.027
	(0.132)	(0.116)	(0.077)	(0.053)	(0.113)	(0.085)
Controls	✓	✓	✓	✓	✓	✓
Firm FEs		✓		✓		✓
N	2201	2201	2147	2147	2147	2147

Notes: This table shows estimated coefficients for equation 1.5. Dependent variables are: the RTI index in columns (1) and (2), the share of Manufacturing Skills in columns (3) and (4), and the share of IT skills in columns (5) and (6). All variables are year averages. Capital and Number of Employees are included as controls. Period: 2013-2014. $Post$ is a dummy taking value 1 in year 2014. Exposure is calculated using job ads posted in 2012. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

1.5.3 Measuring Effects on Trade

Next I turn to difference-in-differences estimates on trade volumes. I regress the model in equation 1.5 on the log of the average monthly imports, the log of average monthly exports and the difference between the two, which can be interpreted as the log of the

[26]Because of the data availability, this placebo exercise differs from the main analysis in terms of sample size. Indeed, the exposure measure is computed using only job advertisements posted in 2012, rather than 2012-2013, and the (placebo) treatment sample is 2014, rather than 2015-2015. To eliminate the concern that the differences between the placebo analysis and the main analysis can be due to the dimensionality of the two subsamples, in table 1.G1 in appendix I replicate the placebo analysis shifting all by one year, to include 2015 as treatment and I show that results are consistent with the ones presented in table 1.B4

import-export ratio. As before, I consider firm-year cells and use data from 2014 as the "pre" period and 2015-2016 as the "post" period. Table 1.7 summarizes the results.

Table 1.7: Trade - Results

	OLS Estimation					
Dependent Variable:	*Log Import*		*Log Export*		*L Imp - L Exp*	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.033	-0.207**	0.117	0.042	-0.084	-0.249
	(0.456)	(0.099)	(0.625)	(0.145)	(0.389)	(0.154)
Exposure (sub)	-1.875**		-2.775**		0.900	
	(0.794)		(1.088)		(0.677)	
Exposure × Post	0.711	0.716***	0.511	0.127	0.200	0.589*
	(0.995)	(0.217)	(1.364)	(0.317)	(0.848)	(0.335)
Controls	✓	✓	✓	✓	✓	✓
Firm FEs		✓		✓		✓
N	4232	4232	4232	4232	4232	4232

Notes: This table shows estimated coefficients for equation 1.5 on log imports in columns (1) and (2). In columns (3) and (4) the dependent variable is log exports and in columns (5) and (6), it is the Import-Export ratio in logs. Dependent variables are log of monthly averages, observed yearly. Capital and Number of Employees are included as controls. $Post$ is a dummy taking value 1 in years 2015 and 2016. Robust standard errors in parentheses. Table 1.G3 in the appendix presents estimates with standard errors clustered at firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Consistent with the correlation patterns reported in table 1.2, the exposure measure is negatively correlated with both imports (column 1) and exports (column 2) before the shock. Looking at the treatment effects, the estimates in column (2) show that the shock increases monthly imports by 0.716 log points for a unitary increase in the exposure measure.[27] This result is in line with the prediction that firms with workers that can be more easily substituted by imported inputs or by machines react strongly to the change in trade prices by importing more goods. This suggests that an increase in imports might be a relevant channel the is the channel through which these firms change the composition of their workforce. A further indicator in support of this thesis is provided by the fact that I do not find any difference between exposed and less exposed firms in terms of exports after the shock (column 4). Therefore, any change in terms of exports between firms can not be attributed to different compositions of their labor

[27] Table 1.G3 in the appendix shows that these results are robust to clustering standard errors at the firm level.

force.

I then estimate equation 1.6 with the logarithm of imports and exports as outcome variables. Figure 1.6 plots the treatment effects ρ_l, and highlights no changes in trade for different levels of exposure before the shock and an increase in imports following the shock.

Finally, I replicate the main analysis separately for transactions between Swiss firms and the Euro area and that made between Swiss firms and the rest of the world and also for transactions denominated in Euro, Chf and US Dollars. As shown in figure 1.H2 in appendix, most of the effect found in table 1.7 is due to an increase in imports, relative to export, from the Euro area and for transactions denominated in Euro.

1.5.4 Trade Driven skill-biased Effects of the Shock

To further investigate the link between changes in trade and in labor demand after the shock, and to combine the two results above, I estimate a two stage least square model on the effects of imports on skill requirements and I use the interaction between the exposure measure and the shock as an instrument for imports. In practice, the first stage is an estimation of the model in equation 1.5 on imports, as in the section above. Equations 1.7 and 1.8 represent the first and second stage models, respectively.

$$import_{it} = \alpha_2 + \gamma_{2,i} + \lambda_{2,t} + \delta_2 X_{it} + \beta_2 sub_i \times D_{t=post} + \varepsilon_{it} \tag{1.7}$$

$$Y_{it} = \alpha_3 + \gamma_{3,i} + \lambda_{3,t} + \delta_3 X_{it} + \beta_3 \widehat{import}_{it} + \varepsilon_{it} \tag{1.8}$$

The coefficient of interest is β_3, which shows the change in the outcome generated by a percentage increase in import. This model is based on the assumption that imports are the channel through which different values of exposure translate into changes in skill requirements. The difference in the interaction coefficients between imports and exports highlighted in table 1.7 supports this assumption. Table 1.8 presents the estimated second stage coefficients together with the estimates of simple models of skill requirements and imports.[28]

As shown in columns (1) and (3), once firm fixed effects are taken into account, the strong correlation between imports and the routine intensity index, or manufacturing

[28]The first stage coefficients are reported in the first two columns of table 1.7.

Figure 1.6: Event Study - Imports and Exports

(a) ES - Import

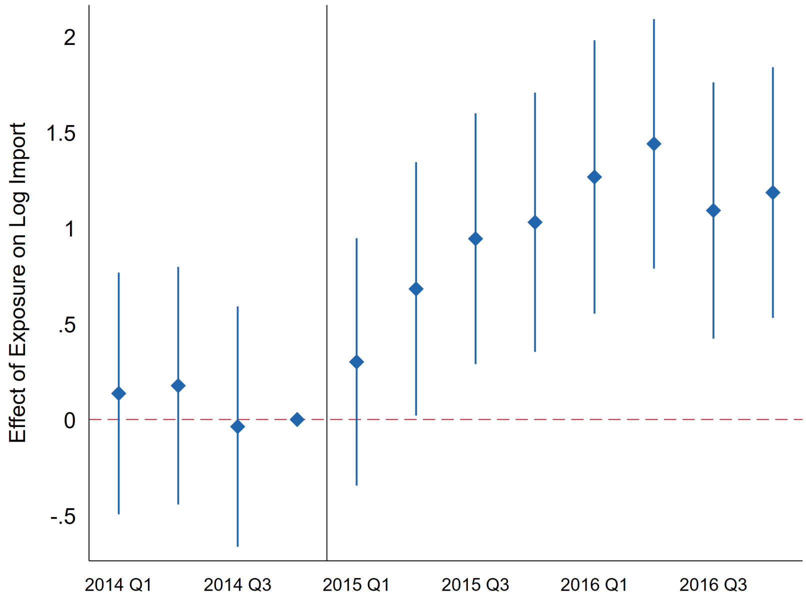

(b) ES - Export

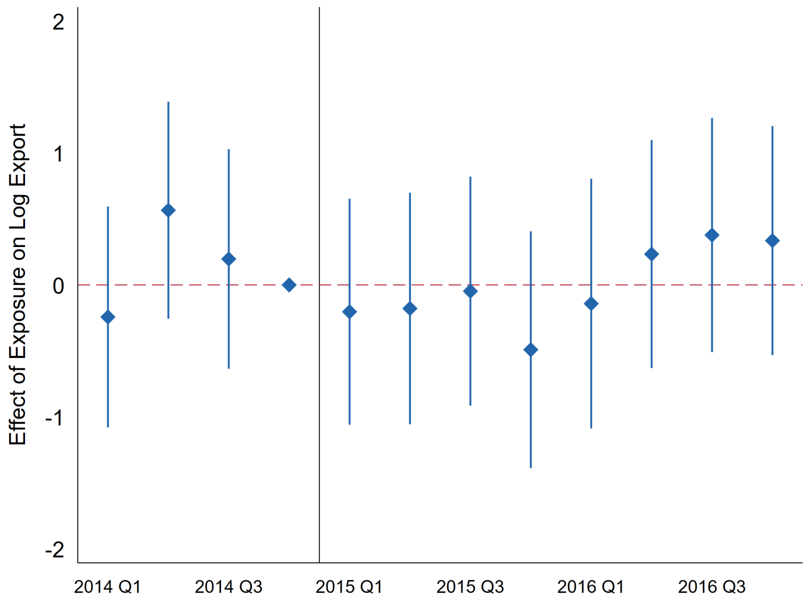

Notes: The figure plots the estimated treatment effects ρ_l in equation 1.6, and their confidence intervals, for Log Imports in panel (a) and Log Exports in panel (b).

Table 1.8: Trade and Skill Requirements - Results

OLS and 2SLS Estimation						
Dependent Variable:	*RTI*		*Manuf. Skills*		*IT Skills*	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
Post	-0.021**	0.010	-0.005	0.006	0.018***	0.015
	(0.008)	(0.018)	(0.004)	(0.009)	(0.007)	(0.011)
Log Import	0.009	-0.274*	-0.003	-0.107*	0.030***	0.062
	(0.009)	(0.155)	(0.003)	(0.065)	(0.006)	(0.081)
Controls	✓	✓	✓	✓	✓	✓
Firm FEs	✓	✓	✓	✓	✓	✓
N	4232	4232	4156	4156	4156	4156
KP F Stat		10.085		10.742		10.742

Notes: This table shows estimated coefficients for equation **??**. Dependent variables are: the RTI index in columns (1) and (2), the share of Manufacturing Skills in columns (3) and (4), and the share of IT skills in columns (5) and (6). All variables are year averages. Capital and Number of Employees are included as controls. Lof Import in columns (2), (4) and (6) is instrumented with Exposure × Post. *Post* is a dummy taking value 1 in years 2015 and 2016.
Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

skills vanishes. Conversely, the share of jobs requiring IT skills significantly increases with imports. Turning to the two stage least squares results, column (1) indicates that, for each additional log-unit in monthly imports after the shock, firms posted jobs with a 20.6% less routine index. Column (2) shows that the demand for manufacturing skills decreases significantly in response to an appreciation driven surge in imports. Workers with manufacturing skills are probably those that suffered the most from the appreciation, as an increase in imports by one log point decreases the share of job postings requiring manufacturing skills by more than 90%.

1.6 CONCLUSION

In this paper, I study how changes in international market prices affect trade and skill requirements for manufacturing firms. To capture exogenous variation in prices, I exploit an unexpected appreciation of the Swiss franc. On January 15, 2015 the Swiss National Bank abandoned a floor to the exchange rate with respect to the euro generating a sudden appreciation of the local currency of about 15%, which remained rather stable in the following period. This unforeseen appreciation immediately reduced the profits of exporting firms while simultaneously created new opportunities for importing firms.

Using transaction-level import and export data together with information on online job postings for more than 1,700 Swiss manufacturing firms, I document three facts. First, I demonstrate that firm imports positively correlate with job postings of non-routine jobs and with IT skills. In contrast, imports negatively affect the demand for manufacturing skills, traditionally linked to the main activity of most firms in the sample. Second, I show that firms that were positively exposed to substitutability reacted to the shock by importing more and by increasing their demand for skilled labor. Specifically, a one standard deviation increase in the exposure to substitutability generates a post-shock increase in monthly imports by 8.3% and in the routine intensity indicator by 1.7%. Third, I quantify the effects of exposure-induced-imports to labor demand. I find that for each additional log monthly import, a firm reduces the routine intensity associated to its labor demand by 0.274 (2/3 of standard deviation) and the share of jobs requiring manufacturing skills by 0.107 (1/2 of standard deviation).

The first implication of this study is that trade triggers skill-biased labor demand, since offshorable and automatable skills are substitutes for imported inputs and capital. The second implication is that exchange rate policy is labor market policy: preventing the local currency from appreciation, vis-a-vis trading partners, slows down the upskilling. To the best of my knowledge, this is the first paper analyzing a sudden and extraordinary change in real prices to provide evidence of short and medium term effects of imports on detailed skill requirements at the firm level.

APPENDIX

1.A THE SHOCK - EXPECTATIONS

Figure 1.A1: The shock - a sudden and permanent variation

(a) KOF Consensus Forecast

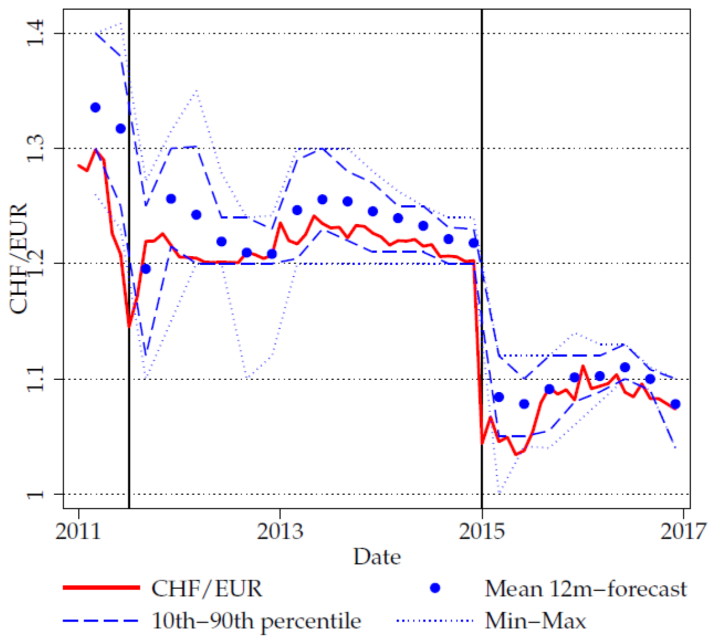

(b) Google Trend Search for "Swiss Franc"

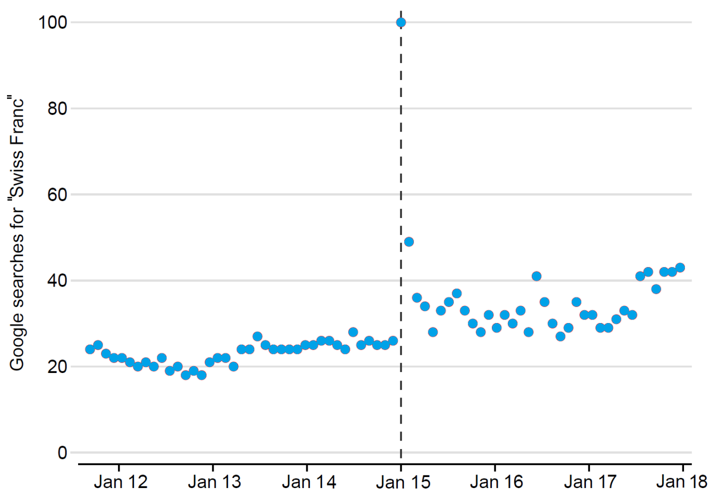

Notes: Panel (a) reports figure 1 (a) in Kaufmann and Renkin (2019). It plots the EUR/CHF exchange rate and its 12 month forecast from the KOF Consensus panel. The figure in panel (b) shows the number of times that the term "*Swiss Franc*" was searched on Google, in percentage with respect to the number of times it was searched in January 2015.

1.B ADDITIONAL DESCRIPTIVE STATISTICS

Table 1.B1: List of Industries

Industry (NACE 2dg)	**N. of firms** (1)	**Share** (2)	**N. of Postings** (3)	**Share** (4)
Manufacture of machinery and equipment n.e.c.	375	.214	12781	.16
Manufacture of fabricated metal products, except machinery and equipment	303	.173	6907	.086
Manufacture of computer, electronic and optical products	208	.119	13373	.167
Manufacture of food products	180	.103	11321	.142
Manufacture of electrical equipment	103	.059	4023	.05
Manufacture of chemicals and chemical products	85	.049	3935	.049
Manufacture of rubber and plastic products	83	.047	2414	.03
Manufacture of wood and of products of wood and cork, except furniture	69	.039	1079	.014
Manufacture of basic pharmaceutical products and pharmaceutical preparations	54	.031	11892	.149
Other manufacturing	51	.029	1612	.02
Manufacture of basic metals	43	.025	1683	.021
Manufacture of furniture	36	.021	851	.011
Repair and installation of machinery and equipment	31	.018	876	.011
Manufacture of textiles	25	.014	981	.012
Manufacture of motor vehicles, trailers etc.. i	23	.013	1015	.013
Manufacture of paper and paper products	21	.012	719	.009
Manufacture of other non metallic mineral products	19	.011	873	.011
Manufacture of beverages	16	.009	832	.01
Manufacture of other transport equipment	13	.007	1923	.024
Printing and reproduction of recorded media	4	.002	120	.002
Manufacture of wearing apparel	4	.002	210	.003
Manufacture of tobacco products	3	.002	294	.004
Manufacture of leather and related products	2	.001	41	.001
Manufacture of coke and refined petroleum products	1	.001	105	.001
Total	1752		79860	

Notes: This table reports the average number of firms and postings by Sector. Period 2014-2016.

Table 1.B2: List of Occupations

Industry (ISCO 2dg)	**N. of Postings** (1)	**Share** (2)	**RTI** (3)	**Offshorability** (4)
Administrative and Commercial Managers	11130	.139	-.9582016	.3443703
Science and Engineering Professionals	7780	.097	-.9525113	.5246673
Business and Administration Professionals	6678	.084	-.9164736	1.1048
Science and Engineering Associate Professionals	6480	.081	-.3026362	.2347695
Metal, Machinery and Related Trades Workers	4409	.055	-.1608545	-.7707986
Business and Administration Associate Professionals	4316	.054	-.4285929	.7969047
Numerical and Material Recording Clerks	3692	.046	.6166951	1.173425
Information and Communications Technology Professionals	3545	.044	-.8101774	1.501084
Missing	3495	.044		
Sales Workers	3414	.043	-.3356114	-.2576303
Electrical and Electronics Trades Workers	2988	.037	-.656814	-1.98177
Production and Specialized Services Managers	2966	.037	-.9292242	.0308005
Other Clerical Support Workers	2602	.033	.5131762	1.535775
Stationary Plant and Machine Operators	1990	.025	.7820621	1.045133
Chief Executives, Senior Officials and Legislators	1723	.022	-.9573376	.5057079
Food Processing, Woodworking, Garment and Related Trades Workers	1216	.015	.2575031	.004385
Assemblers	1197	.015	.9985917	.4333546
Legal, Social and Cultural Professionals	1093	.014	-.6072029	1.126926
General and Keyboard Clerks	1052	.013	.7981312	.9541268
Information and Communications Technicians	1052	.013	-.1757562	.5445752
Customer Services Clerks	1034	.013	-.1149181	1.102903
Health Associate Professionals	910	.011	.1469327	-.4876955
Hospitality, Retail and Other Services Managers	755	.009	-.4204742	-.0660189
Labourers in Mining, Construction, Manufacturing and Transport	729	.009	.1979881	.279027
Teaching Professionals	685	.009	-.9832062	-.3442284
Personal Services Workers	593	.007	-.5209406	.3340856
Building and Related Trades Workers (excluding Electricians)	516	.006	-.8654285	-.5524015
Health Professionals	428	.005	-.8838043	-1.59927
Drivers and Mobile Plant Operators	426	.005	.0308621	-.5969043
Handicraft and Printing workers	325	.004	-.006214	-.3120773
Cleaners and Helpers	219	.003	-.9972603	.9090365
Refuse Workers and Other Elementary Workers	136	.002	-.7892157	.2633111
Legal, Social, Cultural and Related Associate Professionals	133	.002	-.4523416	.7680697
Protective Services Workers	102	.001	-1	-.8487132
Food Preparation Assistants	20	0	-.9777778	.9009433
Market-oriented Skilled Agricultural Workers	13	0	-.5780219	-.1445741
Personal Care Workers	12	0	-.7916667	-.7419518
Agricultural, Forestry and Fishery Labourers	4	0	-.5384616	.751695
Market-oriented Skilled Forestry, Fishery and Hunting Workers	2	0	-.5641026	-.3713484
Total	79860			

Notes: This table reports the average number of postings by Occupation, together with the RTI and the offshorability index attributed to the job. Period 2014-2016.

Table 1.B3: Summary statistics - Skill Families

Variable	**Mean** (1)	**Std. Dev.** (2)	**Min.** (3)	**Max.** (4)	**N** (5)
Higher Education	0.087	0.282	0	1	79860
RTI	-0.487	0.62	-1	1	72943
Job Offshorability	0.391	1.025	-3.011	2.661	72055
Administration	0.025	0.155	0	1	79860
Analysis	0.037	0.188	0	1	79860
Business	0.42	0.494	0	1	79860
Design	0.086	0.281	0	1	79860
Economics and Social	0.015	0.121	0	1	79860
Education and Training	0.031	0.172	0	1	79860
Energy and Utilities	0.021	0.143	0	1	79860
Engineering	0.092	0.289	0	1	79860
Finance	0.173	0.378	0	1	79860
Health Care	0.054	0.226	0	1	79860
Human Resources	0.072	0.259	0	1	79860
Industry Knowledge	0.512	0.5	0	1	79860
Information Technology	0.467	0.499	0	1	79860
Manufacturing	0.082	0.274	0	1	79860
Marketing and PR	0.062	0.241	0	1	79860
Media and Writing	0.028	0.164	0	1	79860
Personal Care, Services	0.035	0.185	0	1	79860
Sales	0.13	0.336	0	1	79860
Science and Research	0.09	0.286	0	1	79860
Logistics	0.059	0.235	0	1	79860

Notes: This table reports summary statistics for vacancies posted in the period 2014-2016. Each observation is a job advertisement.

Table 1.B4: Manufactruring Skills

	Manufacturing Skills	
Skill Cluster	**Share (percent)**	**Cumulative (percent)**
	(1)	(2)
Metal Fabrication	47.61	47.61
Product Development	20.77	68.38
Manufacturing Processes	11.89	80.27
Materials Science	6.15	86.42
Materials Process	6.07	92.49
Welding	4.06	96.55
Computer-Aided Manufacturing	1.16	97.71
Machinery	0.98	98.69
Machine Tools	0.41	99.11
Manufacturing Design	0.37	99.47

Notes: This table shows the top 10 skill clusters in the skill family "Manufacturing Skills". For each cluster the share of skills with that cluster over all skills in the family is reported. Sample: 2014-2016.

Table 1.B5: IT Skills

	IT Skills	
Skill Cluster	**Share (percent)**	**Cumulative (percent)**
	(1)	(2)
Microsoft Office and Productivity Tools	13.09	13.09
Other Programming Languages	10.57	23.66
Enterprise Resource Planning (ERP)	10.54	34.21
Software Development Tools	8.39	42.6
Technical Support	7.3	49.89
Scripting Languages	6.81	56.7
IT Management	6.71	63.41
Advanced Microsoft Excel	5	68.41
Internet Protocols	3.7	72.11
Computer Hardware	3.09	75.2
Operating Systems	2.13	77.33
Database Management Systems	1.62	78.95
Application Development	1.56	80.51
Microsoft Development Tools	1.45	81.95
Data Management	1.31	83.26
Microsoft Windows	1.17	84.42
Internet Services	1.08	85.51
Network Configuration	0.95	86.46
Database Administration	0.9	87.36
Internet Security	0.84	88.2

Notes: This table shows the top 20 skill clusters in the skill family "Information Technology Skills". For each cluster the share of skills with that cluster over all skills in the family is reported. Sample: 2014-2016.

1.C Examples - Job Postings

Figure 1.C1: Examples - Skills

(a) Example 1

```
In [8]: df['title'][0]
Out[8]: "Teamleiter Logistik'"

In [9]: df['content'][0]
Out[9]: "Aufgaben:  Fuhren des Teams Logistik Administration (4 Personen) Disposition und Organisation von Transporten mit externen Die
        nstleitern im In- und Ausland Erstellen von samtlichen Versanddokumenten fur internationale Lieferungen Akkreditivabwicklung so
        wie Erstellung und Uberwachung der Dokumente Kontakt zu Spediteuren, Banken, Versicherungen, Zollstellen usw. Stammdatenpflege
        im ERP System  Anforderungen:  Abgeschlossene kaufmannische Ausbildung mit mehrjahriger Berufserfahrung Weiterbildung im Bereic
        h Aussenhandel und oder Logistik/Spedition von Vorteil Fuhrungserfahrung Sehr gute Englisch- und MS Office-Kenntnisse Motiviert
        e, emphatische, aufgestellte und dienstleistungsorientierte Personlichkeit  '"

In [10]: df['skills_ger'][0]
Out[10]: "[' Logistik' , ' Produktionsanlage' , ' Schreiben' , ' Fuhrung' , ' MS Office' , ' ERP' ]"

In [11]: df['skills_eng'][0]
Out[11]: "['Logistic' , ' Production plant' , ' Writing' , ' Management' , ' Office applications' , ' Enterprise Resource Planning (ER
         P)' , ]'"

In [12]: df['skill_family'][0]
Out[12]: "[' Logistic' , ' Sales' , ' Media and Writing' , ' Business' , ' Information Technology' , ' Information Technology' ,]'"
```

(b) Example 2

```
In [14]: df['title'][1]
Out[14]: 'Industrielackierer (m/w)'

In [15]: df['content'][1]
Out[15]: "Industrielackierer (m/w) / Temporar (100%) ab Sofort   Anstellungsart Temporar Region St.Gallen/ Appenzell  Fur unseren Kunde
         n, ein Maschinenbau-Unternehmen aus der Region Uzwil, suchen wir per sofort oder nach Vereinbarung einen Industrielackierer. Ih
         r Aufgabengebiet: Aufgaben:-Lackieren hochkomplexer Teile (2K- Nasslack)-Vorbehandlung/Vorbereitung -Vorbereiten und Erstellen
         der Lacke-Prufung der Oberflachenqualitat Ihr Profil -Abgeschlossene Berufsausbildung als Industrie- oder Autolackierer-Gute Ke
         nntnisse der Objektvorbereitung-Breites Wissen uber Lacksysteme-Ausgepragte Fahigkeit im farblichen Sehen-Teamfahigkeit -Selbst
         andige, flexible Arbeitsweise-Deutsch in Wort und Schrift-Gute korperliche Verfassung-Bereitschaft fur Schichtarbeit (2-Schich
         t) FilialeFiliale Weinfelden  Ihr AnsprechpartnerRoland Nageli'"

In [16]: df['skills_ger'][1]
Out[16]: "[' CNC' , ' Mitarbeitermotivation' ,  ' Frasen' , ' CAM' ]"

In [17]: df['skills_eng'][1]
Out[17]: "[' CNC' , ' Employee motivation' , ' Milling' , ' CAM' ]"

In [18]: df['skill_family'][1]
Out[18]: "[' Manufacturing and Production' , ' ---' ,  ' Manufacturing and Production' , ' Design' ]"
```

Notes: This figure shows the job descriptions of two job postings, the skills extracted from it, their translation and the associated skill families. Example (a) refers to a *Team leader for logistic* while example (b) refers to an *Industrial Painter*.

1.D Distributions

1.D.1 Trade

Figure 1.D1: Distribution of Imports and Exports

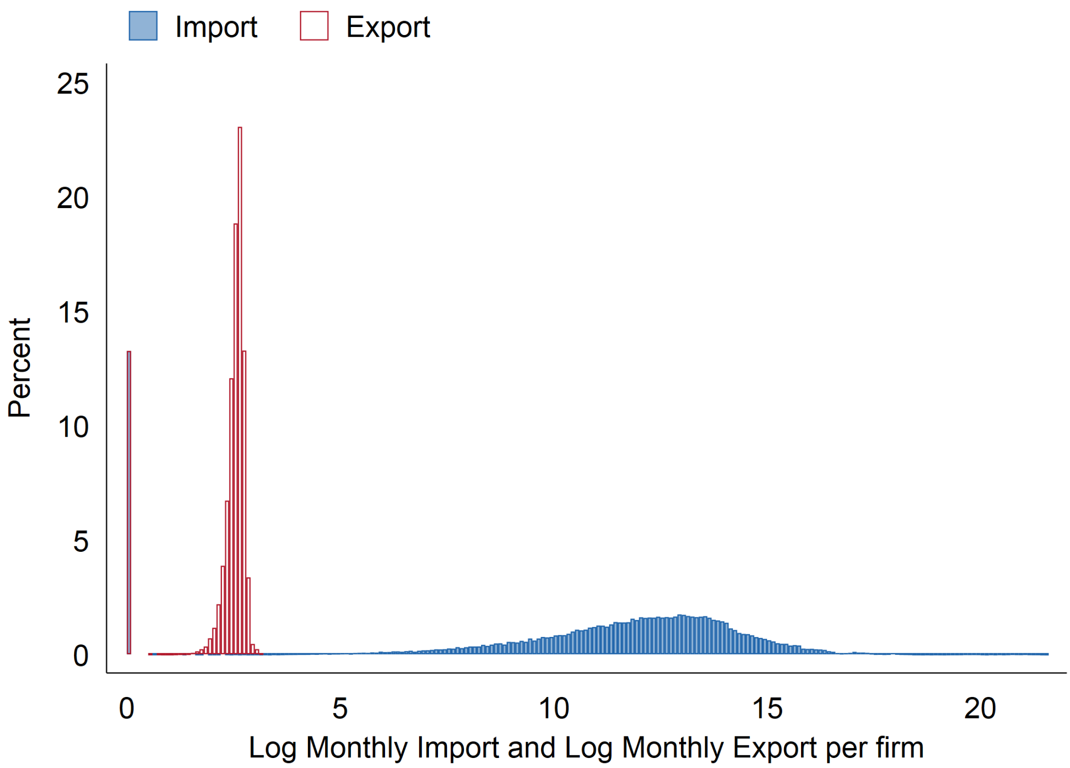

Notes: This figure shows the the distribution of log imports and log exports for the period 2014-2015. Each observation is a firm-month cell.

1.D.2 Skills and Skill Families

Figure 1.D2: Distribution of Skills

(a) Skills (b) Skill Families

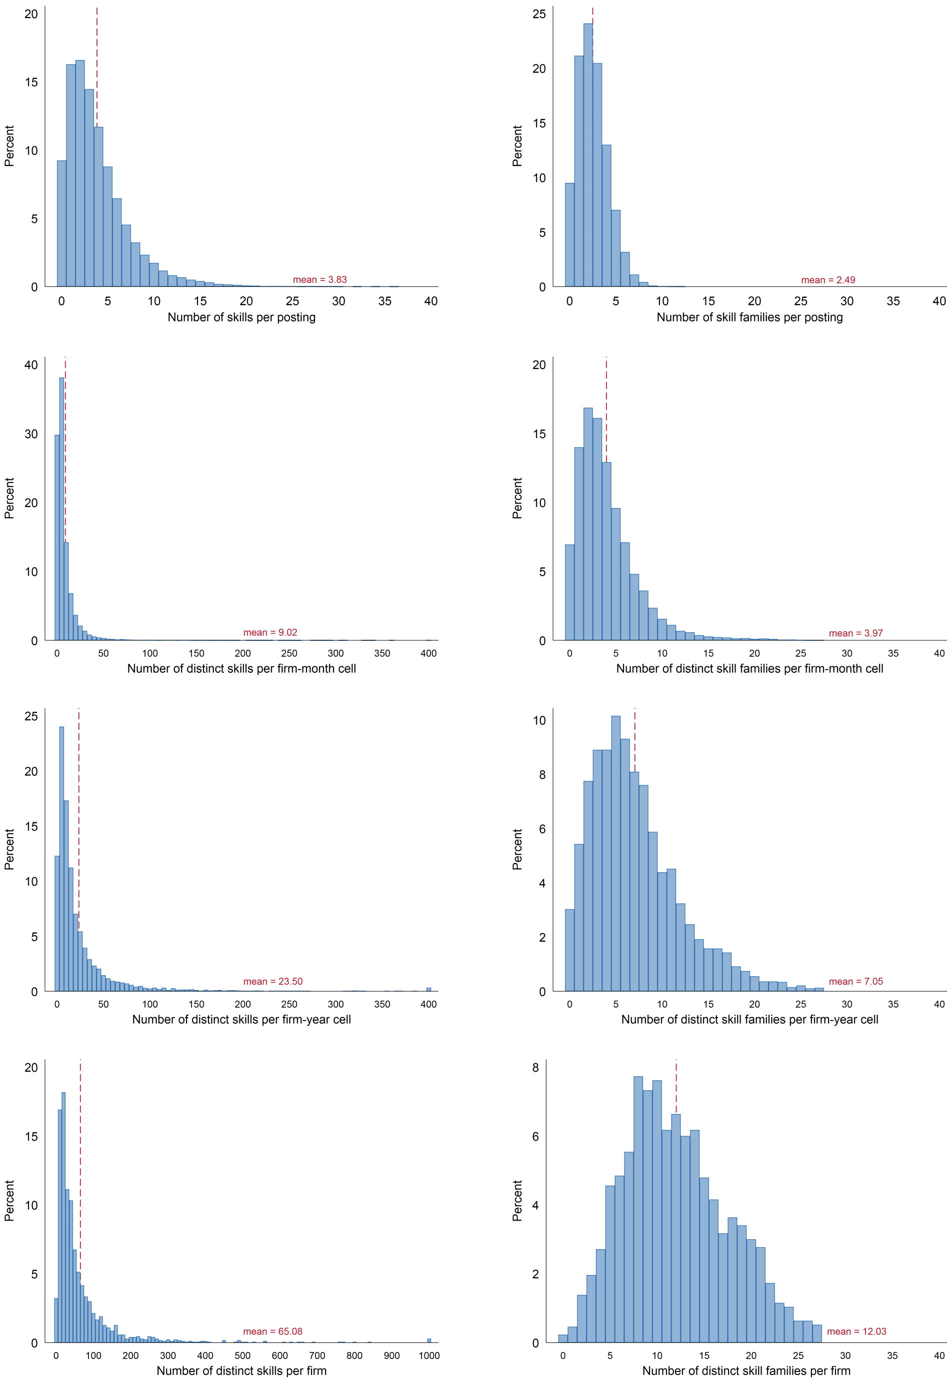

Notes: This figure shows the distribution of skills and skill families by job, firm-month cell, firm-year cell, and firm for the period 2014-2016

1.D.3 RTI, OFFSHORABILITY AND SUBSTITUTABILITY

Figure 1.D3: Distributions

(a) Routine Intensity Index

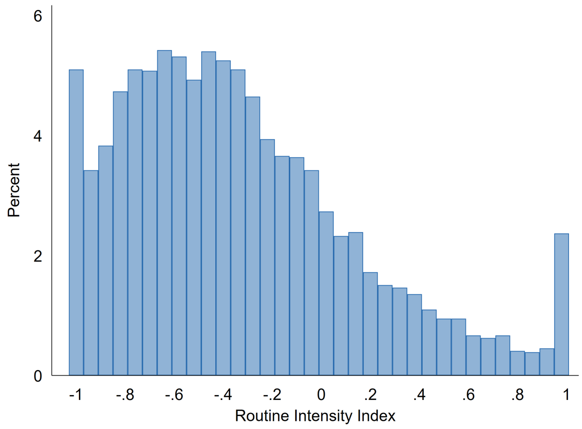

(b) Offshorability Index

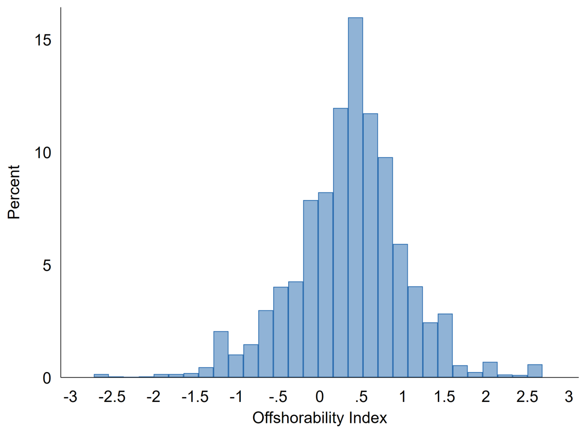

(c) Substitutability Index

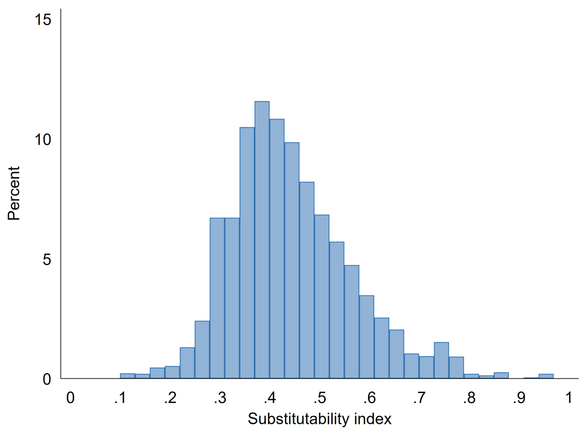

Notes: This figure shows the distribution of the routine intensity indicator (RTI), the offshorability index and the substitutability index for the period 2014-2015. Each observation is a firm-year cell.

1.E Skills, Occupations and Industries

Figure 1.E1: Skills and Occupations

(a) Share of Postings by Occupation within Skill Requirements

(b) Share of Postings by Skill Requirement within Occupations

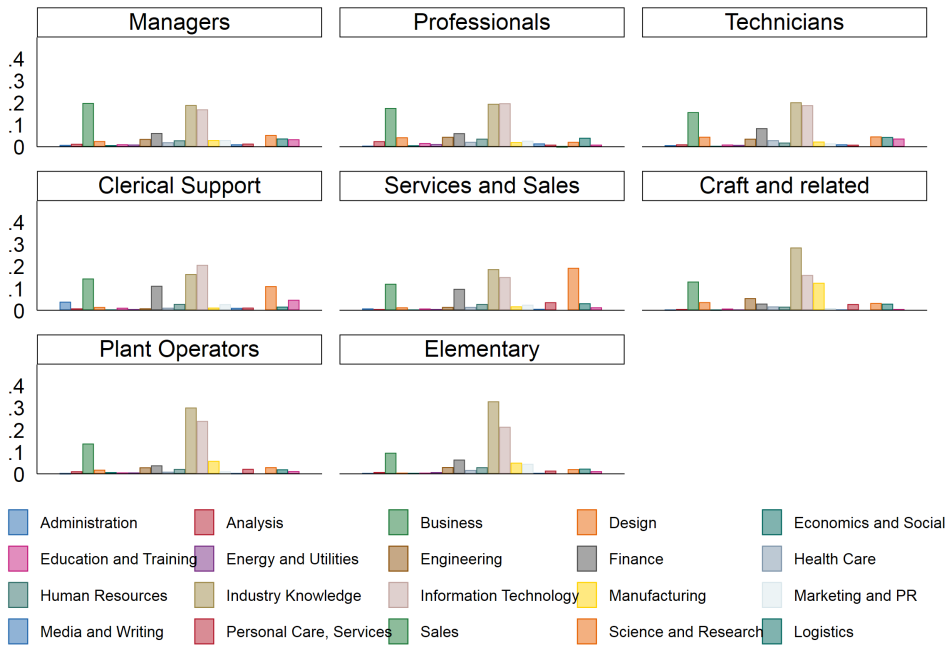

Notes: This figure shows the heterogeneity between skills and occupations. Panel A shows the share of postings by occupation within skill requirements. Panel B shows the share of Job Postings by skill requirement within each occupation. Each bar indicates the portion of job postings in an occupation and with a given skill requirement over the total number of job postings with the same skill requirement in Panel A, and over the total number of job postings in that occupation in Panel B.

Figure 1.E2: Skills and Industries

(a) Share of Postings by Sectors within Skill Requirements

(b) Share of Postings by Skill Requirement within Sectors

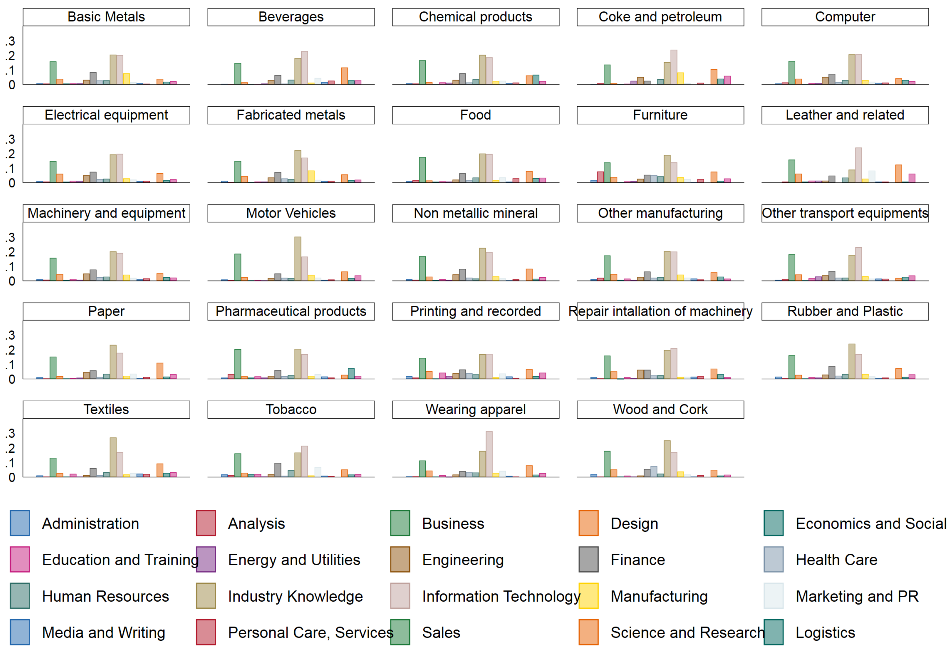

Notes: This figure shows the heterogeneity between skills and sub manufacturing sectors. Panel A shows the share of postings by sector within skill requirements. Panel B shows the share of job postings by skill requirement within each sector. Each bar indicates the portion of job postings in an occupation and with a given skill requirement over the total number of job postings with the same skill requirement in Panel A, and over the total number of Job Postings in that sector in Panel B.

1.F Pre-Shock: Correlations and Distributions

1.F.1 Linear Correlations

Table 1.F1: Correlation Coefficients

Variable	N ads (1)	High Educ (2)	RTI (3)	Offshorability (4)	N Empl (5)	Capital	Op Rev	Import	Export
Panel A: *Job Postings*									
Number of ads	1.000								
Bachelor Degree or Higher	0.152	1.000							
RTI	-0.120	-0.210	1.000						
Job Offshorability	0.026	0.115	0.131	1.000					
Panel B: *Firm*									
Number of Employees	0.604	0.096	-0.043	0.008	1.000				
Capital	0.388	0.100	-0.055	0.017	0.679	1.000			
Operating Revenue	0.532	0.116	-0.053	0.010	0.993	0.671	1.000		
Panel C: *Trade*									
Total Import	0.496	0.085	-0.046	0.000	0.707	0.446	0.737	1.000	
Total Export	0.711	0.119	-0.057	0.006	0.888	0.530	0.806	0.888	1.000

Notes: This table reports correlation coefficients between the main variables in the analysis for the year 2014. Each observation corresponds to a firm.

1.F.2 TRADE BY SECTOR

Table 1.F2: Trade by Sector

Sector (NACE 2dg)	**Import (in Mln CHF)** (1)	**Export (in Mln CHF)** (2)
Manufacture of basic pharmaceutical products and pharmaceutical preparations	453.82	1054.23
Manufacture of basic metals	484.95	391.86
Manufacture of other transport equipment	56.62	129.12
Manufacture of coke and refined petroleum products	49.63	102.72
Manufacture of chemicals and chemical products	24.78	43.42
Manufacture of computer, electronic and optical products	18.38	42.05
Manufacture of paper and paper products	20.44	29.43
Manufacture of machinery and equipment n.e.c.	9.96	26.88
Manufacture of motor vehicles, trailers, etc...	17.32	25.89
Manufacture of other non metallic mineral products	18.28	24.81
Manufacture of electrical equipment	11.90	21.92
Manufacture of textiles	11.34	21.03
Other manufacturing	5.98	20.94
Manufacture of food products	15.51	19.65
Manufacture of wearing apparel	12.34	13.15
Manufacture of leather and related products	5.42	11.18
Manufacture of rubber and plastic products	8.01	10.70
Manufacture of fabricated metal products, except machinery and equipment	4.52	10.44
Manufacture of wood and of products of wood and cork, except furniture	4.97	5.72
Printing and reproduction of recorded media	10.53	5.18
Manufacture of furniture	4.12	4.41
Manufacture of beverages	18.11	2.71
Repair and installation of machinery and equipment	1.84	2.31
Manufacture of tobacco products	9.47	1.27

Notes: This table reports the firm average yearly imports (in million CHF) and exports (in million CHF) per sector in 2014.

1.F.3 Trade and Skills

Table 1.F3: Trade and Skill Requirements - Cross Section Correlations in 2014

	OLS Estimation							
Dependent Variable:	*Log Import*				*Log Export*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
N ads	0.014***			0.011***	0.017***			0.013***
	(0.003)			(0.003)	(0.004)			(0.003)
RTI		-1.205***		-0.655**		-2.015***		-1.175***
		(0.281)		(0.286)		(0.376)		(0.387)
Administration			-0.748	-0.059			0.016	1.177
			(1.106)	(1.051)			(1.368)	(1.371)
Analysis			1.105	0.509			1.204	0.316
			(0.792)	(0.766)			(1.315)	(1.269)
Business			1.733***	1.421***			2.179***	1.816***
			(0.408)	(0.408)			(0.551)	(0.553)
Design			-1.610**	-1.620**			-0.347	-0.403
			(0.666)	(0.660)			(0.851)	(0.846)
Economics and Social			3.170	2.127			5.684**	4.497
			(2.022)	(1.996)			(2.845)	(2.755)
Education and Training			0.558	0.425			-1.657	-1.352
			(1.145)	(1.149)			(1.754)	(1.628)
Energy and Utilities			-0.010	-0.798			0.967	-0.252
			(2.240)	(2.156)			(2.809)	(2.663)
Engineering			0.210	0.070			1.201	1.026
			(0.676)	(0.689)			(0.905)	(0.917)
Finance			1.632***	1.687***			2.958***	3.181***
			(0.494)	(0.502)			(0.646)	(0.644)
Health Care			-0.650	-0.694			-0.867	-0.937
			(0.784)	(0.784)			(1.005)	(0.993)
Human Resources			0.093	-0.172			0.258	-0.128
			(0.925)	(0.910)			(1.146)	(1.127)
Industry Knowledge			1.084***	0.909**			1.854***	1.670***
			(0.375)	(0.372)			(0.489)	(0.487)
Information Technology			1.085***	0.988***			2.144***	1.993***
			(0.368)	(0.365)			(0.495)	(0.490)
Manufacturing			-1.043**	-0.937*			-0.720	-0.510
			(0.510)	(0.519)			(0.709)	(0.722)
Marketing and PR			-1.063	-1.150			-1.878	-2.264*
			(0.950)	(0.855)			(1.329)	(1.323)
Media and Writing			-0.441	-0.577			-0.564	-0.753
			(1.266)	(1.251)			(1.685)	(1.675)
Personal Care, Services			-4.326***	-4.479***			-8.715***	-8.536***
			(1.622)	(1.671)			(2.016)	(2.067)
Sales			0.013	0.008			-1.085	-0.933
			(0.495)	(0.493)			(0.666)	(0.663)
Science and Research			1.421	1.139			3.168**	2.723**
			(0.909)	(0.916)			(1.333)	(1.342)
Logistics			2.222**	1.961**			1.749	1.501
			(0.922)	(0.902)			(1.399)	(1.331)
Constant	13.064***	12.983***	11.911***	11.789***	12.213***	11.936***	10.031***	9.722***
	(0.123)	(0.142)	(0.292)	(0.293)	(0.159)	(0.190)	(0.362)	(0.370)
Observations	1670	1653	1670	1653	1670	1653	1670	1653

Notes: This table shows the regression coefficients of a cross-sectional linear model for 2014 projecting trade over several measures of skill requirements. Each observation refers to a firm in the sample. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.G Robustness

1.G.1 Main results - restricted sample

Table 1.G1: Skill Requirements - 2014-2015

	OLS Estimation					
Dependent Variable:	*RTI*		*Manuf. Skills*		*IT Skills*	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.030	0.072**	0.029	0.060**	-0.008	-0.022
	(0.067)	(0.032)	(0.034)	(0.025)	(0.051)	(0.038)
Exposure (sub)	1.059***		0.073		-0.052	
	(0.101)		(0.051)		(0.076)	
Exposure × Post	-0.130	-0.172**	-0.098	-0.150***	0.054	0.083
	(0.146)	(0.071)	(0.074)	(0.055)	(0.112)	(0.082)
Controls	✓	✓	✓	✓	✓	✓
Firm FEs		✓		✓		✓
N	2201	2201	2147	2147	2147	2147

Notes: This table shows estimated coefficients for equation 1.5. Dependent variables are: the RTI index in columns (1) and (2), the share of Manufacturing Skills in columns (3) and (4), and the share of IT skills in columns (5) and (6). All variables are year averages. Capital and Number of Employees are included as controls. Period: 2014-2015. $Post$ is a dummy taking value 1 in year 2015. Exposure is calculated using job ads posted in 2013.
Robust standard errors in parentheses. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

1.G.2 MAIN RESULTS - CLUSTERING

Table 1.G2: Skill Requirements - Clustering

	OLS Estimation					
Dependent Variable:	*RTI*		*Manuf. Skills*		*IT Skills*	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.035	0.067	0.015	0.032	-0.019	-0.011
	(0.051)	(0.050)	(0.026)	(0.025)	(0.035)	(0.034)
Exposure (sub)	1.059***		0.073		-0.052	
	(0.109)		(0.055)		(0.081)	
Exposure × Post	-0.175	-0.196*	-0.056	-0.086*	0.090	0.069
	(0.115)	(0.114)	(0.050)	(0.049)	(0.077)	(0.075)
Firm FEs		✓		✓		✓
Cluster Firm ID	✓	✓	✓	✓	✓	✓
N	4232	4232	4121	4121	4121	4121

Notes: This table shows estimated coefficients for equation 1.5. Dependent variables are: the RTi index in columns (1) and (2), the share of Manufacturing Skills in columns (3) and (4), and the share of IT skills in columns (5) and (6). All variables are year averages. Capital and Number of Employees are included as controls. $Post$ is a dummy taking value 1 in years 2015 and 2016.
Standard Errors Clustered at firm level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.G3: Trade - Clustering

	OLS Estimation					
Dependent Variable:	*Log Import*		*Log Export*		*L Imp - L Exp*	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.033	-0.207	0.117	0.042	-0.084	-0.249
	(0.225)	(0.133)	(0.264)	(0.158)	(0.199)	(0.168)
Exposure (sub)	-1.875**		-2.775**		0.900	
	(0.848)		(1.133)		(0.729)	
Exposure × Post	0.711	0.716**	0.511	0.127	0.200	0.589
	(0.514)	(0.306)	(0.585)	(0.346)	(0.443)	(0.383)
Firm FEs		✓		✓		✓
Cluster Firm ID	✓	✓	✓	✓	✓	✓
N	4232	4232	4232	4232	4232	4232

Notes: This table shows estimated coefficients for equation 1.5 on log imports in columns (1) and (2). In columns (3) and (4) the dependent variable is log exports and in columns (5) and (6), it is the Import-Export ratio in logs. Dependent variables are log of monthly averages, observed yearly. Capital and Number of Employees are included as controls. $Post$ is a dummy taking value 1 in years 2015 and 2016.
Standard Errors Clustered at firm level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.H Additional Results

1.H.1 Extensive Margin: Number of Postings

Figure 1.H1: Effect on Number of Postings

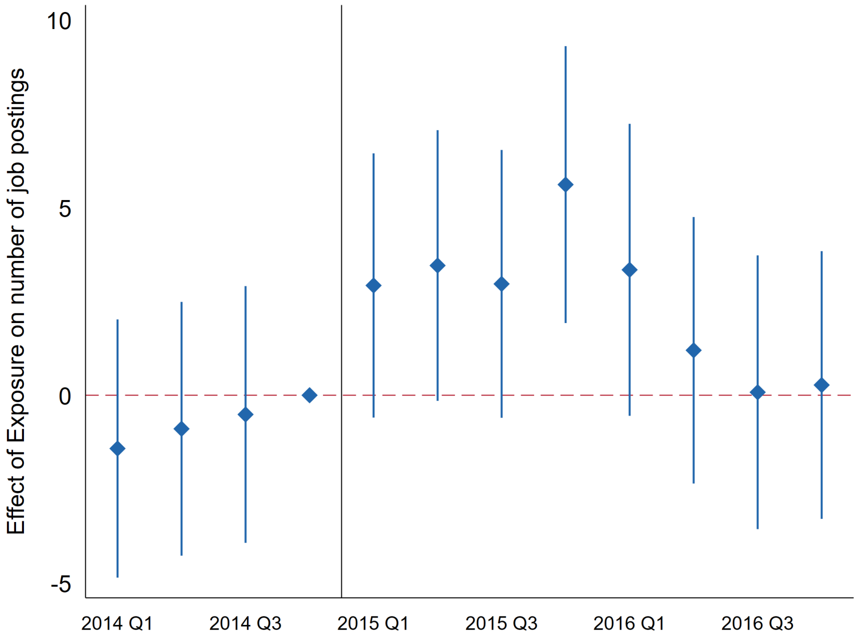

Notes: The figure plots the estimated treatment effects ρ_l in equation 1.6, and their confidence intervals, for a the average number of new job postings per quarter.

1.H.2 Heterogeneous Effects

Figure 1.H2: Heterogeneous Effects

(a) Euro Area vs RoW

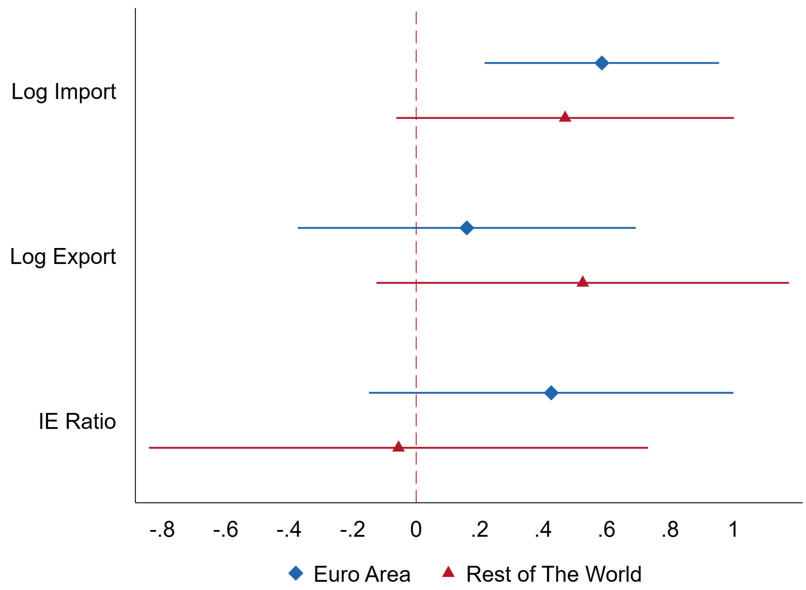

(b) Currencies

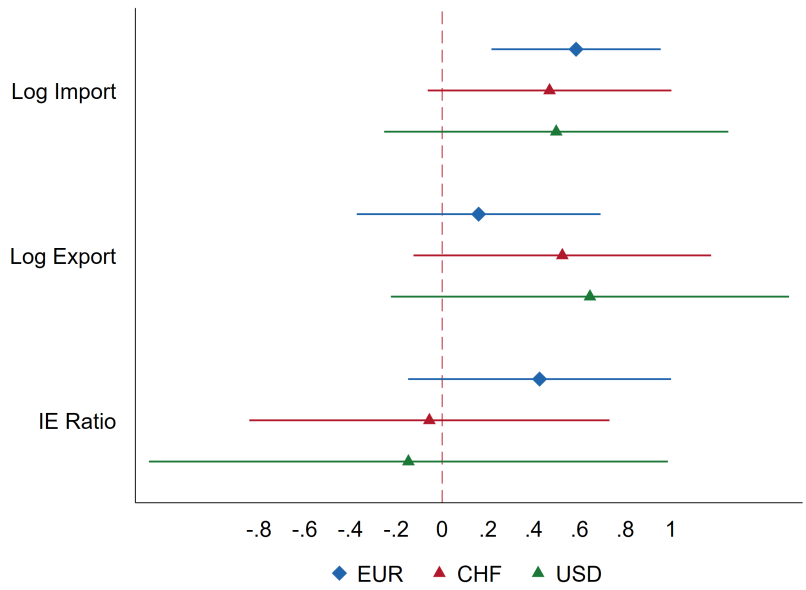

Notes: This figure plots the estimated treatment effects β_1 in equation 1.5, and their confidence intervals, for three outputs, *Log Import*, *Log Exporto*, *IE Ration*, separately for different sets of transactions. Panel (a) reports transactions between Switzerland and the Euro Area versus transactions between Switzerland and the rest of the world. Panle (b) reports separately transactions denominated in Euro, Chf and US Dollars.

CHAPTER 2

GENDER PREFERENCES IN JOB VACANCIES AND WORKPLACE GENDER DIVERSITY[1]

2.1 INTRODUCTION

Rules and informal prohibitions governing gender roles in the workforce have gradually diminished over the last century (e.g., Goldin, 2014). Nevertheless, men and women still tend to work at different firms.[2] How much of the remaining gender segregation in a modern labor market is due to skill or preference differences between males and females versus discriminatory employer preferences is hard to discern, particularly when employers are explicitly forbidden from expressing their preferences (Kuhn and Shen, 2013). Evidence from audit studies (e.g., Booth and Leigh, 2010, Kline et al. (2021)) suggests that employer preferences tend to reinforce gender stereotypes and segregation.[3] As noted by Cahuc et al. (2019), however, it may be difficult to infer the impacts of discriminatory preferences on hiring outcomes from the data that are collected in typical audit studies.

[1]This chapter is coauthored with David Card from the University of California, Berkeley and Rafael Lalive from the University of Lausanne.

[2]Hellerstein et al. (2008) report that a typical female employee of a larger establishment in the U.S. has about 60% female coworkers, while a typical male has only 40% female coworkers. For a broader sample of workers in the U.K., Mumford and Smith (2009), online Appendix Table A2, report corresponding rates of 70% female coworkers for women and 34% for men. Slightly larger gender gaps in exposure to female coworkers are estimated by Card et al. (2016) for Portugal and by Gerard et al. (2018) for Brazil.

[3]Azmat and Petrongolo (2014) and Baert (2017) present overviews of these studies. A number of authors have argued that the tendency to favor the gender group that stereotypically matches the job opening may arise from firms' intentions to reduce search time and improve expected outcomes, rather than from animus discrimination – see e.g., the discussion in Kuhn and Shen (2013).

While the use of stated gender preferences (SGP's) in job recruiting was eliminated in the U.S. in the 1970s, they were still prevalent throughout Europe at the turn of the 21st century, including in Austria, where about 40% of job vacancies listed a preferred gender, despite a 1985 law banning SGP's. In mid 2004, an amendment to the Austrian Equal Treatment Act introduced financial penalties for including gender preferences in job ads, and in early 2005 the *Ombud for Equal Treatment* (OET) agency launched a campaign to inform the public about the law and these penalties. We use the changes after the information campaign to measure the effects of eliminating discriminatory signals at the earliest stage of the job matching process. Vacancy postings from the largest online job board in Austria contain explicit and unambiguous information on stated gender preference, and we combine these data with administrative data on the firms posting the vacancies and the recruits who filled the vacancies. The resulting data set allows us to document how the use of gender preferences in the pre-campaign period varied across employers, and how actual hiring outcomes varied with the SGP's in the associated vacancies.

We show that prior to the campaign, firms with a mainly male workforce were likely to state a preference for males, and vice versa for firms with a mainly female workforce ("stereotypical" preferences). Nevertheless, a small fraction of employers were using SGP's to look for employees that were the opposite gender of the majority of their workforce – posting what we define as "non-stereotypical" vacancies. Vacancies with a gender preference were very likely (>90%) to be filled by a candidate of the preferred gender, even in the case of non-stereotypical preferences. Vacancies with gender preferences were also filled somewhat faster (especially those with non-stereotypical preferences), suggesting that SGP's provided highly informative signals when they were allowed. But would the elimination of these signals change actual hiring outcomes? Or would firms reach the same hiring decisions even after the elimination of SGP's?[4]

To answer these questions we use data on occupation, industry and gender composition of the firm's workplace to classify vacancies in the pre-campaign period into three groups: likely to specify a SGP for males, likely to specify a SGP for females, or likely to specify no preference. We then use the classification model to "tag" vacancies

[4]A similar question arises for recent policies that forbid the collection of data on criminal records or credit histories at the job application phase (Agan and Starr, 2018; Doleac and Hansen, 2020; Bos et al., 2018; Ballance et al., 2020) or the collection of information on salary history (Agan et al., 2020).

before and after the campaign and conduct simple event studies of hiring outcomes for different types of vacancies.

We find that the elimination of firms' abilities to advertise their gender preferences led to a significant rise (+2.5 percentage points) in the fraction of women hired to fill vacancies with a predicted male preference, and a smaller but still significant rise in the fraction of men hired to fill vacancies with a predicted female preference (+1 ppt). The vast majority of vacancies with predicted preferences were stereotypical, (e.g., concordant with the existing gender composition of the firm's workplace), and we show that hires after the campaign increased firm-based gender diversity. These findings suggest that some employers that were posting SGP's prior to the campaign were willing to hire workers of the opposite gender, and actually did so once the composition of their applicant pool was broadened by the elimination of SGP's.

To gain further insights we classify vacancies by the interaction between predicted gender preference and the gender composition of the firm's workplace. This allows us to compare stereotypical versus non-stereotypical vacancies. We find an *increase* in the hiring of women to fill vacancies with both stereotypical and non-stereotypical male predicted preferences (consistent with the objectives of the OET campaign), but a *decrease* in the hiring of women to fill jobs that were predicted to state a non-stereotypical preference for females. The latter impact was presumably an unintended consequence of the campaign, but is interpretable through the lens of a simple model where job seekers use a combination of general information about a firm (like its industry and gender composition), together with information in the SGP, in deciding where to apply. In such a model non-stereotypical SGP's can inform job searchers about preferences that are otherwise unexpected, allowing them to more easily find rare opportunities.

We also examine the impacts of the campaign on three key outcomes for filled vacancies: the time required to fill the vacancy; the wage on the newly established job; and the duration of the new job. We find that the average times required to fill vacancies with a predicted male preference (either stereotypical or non-stereotypical) were unaffected by the campaign, as was the time to fill predicted female vacancies at mainly female workplaces. Thus, for the vast majority of vacancies, matching speed remained constant, though we do find a significant rise in the time to fill vacancies with non-stereotypical female preferences. For wages we find small effects across the board once we control for compositional effects attributable to the gender gap in pay.

For job durations we also find small effects, apart from a decline in the duration of jobs created from vacancies with a stereotypical female gender preference. We attribute the latter impact to a pre-existing trend toward shorter job durations for women in majority female workplaces, rather than to the OET's campaign.

Our findings contribute to the large literature on gender segregation at the workplace level, including Blau (1977), Groshen (1991), Petersen and Morgan (1995), Bayard et al. (2003), and Card et al. (2016). Specifically, we show the importance of direct employer signals – now outlawed in many higher-income countries but still in use in many lower-income countries – in reinforcing gender segregation.

Our findings also contribute to a growing literature concerned with policies that restrict the collection of information at early stages of the job matching process, including criminal records (Agan and Starr, 2018; Doleac and Hansen, 2020), credit histories (Bos et al., 2018; Ballance et al., 2020), and drug use (Wozniak, 2015). As in this literature, the key question in our paper is whether the prohibition of early-stage information affects ultimate hiring outcomes. Consistent with at least some of these studies, we find that it does. We also show that the quality of job matches does not seem to change much after early-stage information is disallowed.

Our work is closely related to studies of gender preferences in Chinese job matching markets, including Kuhn and Shen (2013), Helleseter et al. (2016), and Kuhn et al. (2020). Kuhn and Shen (2013) develop a signaling model of the decision by employers to state a gender preference that compares the cost of screening extra job applications to the expected benefits of being able to evaluate applicants of both genders. Implicit in this model is the assumption – validated by Kuhn et al. (2020) using employer call-backs to applicants – that SGP's are highly predictive of firm's preferences over candidates. We contribute to this literature by observing actual *hiring decisions* (rather than call-backs), by distinguishing between SGP's used by firms with higher and lower shares of female employees, and by studying wages and job durations for filled vacancies. Using hiring outcomes we also show how the elimination of stated gender preferences in job openings led to increases in gender diversity at hiring workplaces.

Most directly, our paper is related to Kuhn and Shen (2021), who study the effects of a decision by the job board in one Chinese city to eliminate SGP's from all posted vacancies. This instantaneous change allows Kuhn and Shen to compare applications to the same vacancy before and after the removal of SGP's, providing a credible design for measuring the reactions of applicants to SGP's. Consistent with our results, they find that the removal of SGP's leads to an increase in the number of applications from the non-preferred gender group. They also show that the share of call-backs to applicants from the non-preferred gender rises after the campaign – again, consistent with our findings on hiring outcomes. We view our results as highly complementary to those of Kuhn and Shen (2021): they have information on applications and employer call-backs, whereas we have information on hiring and job outcomes, allowing us to assess changes in matching efficiency. The consistency of the findings across the two settings and between applications/call-backs versus hiring/job outcomes is highly reassuring, and suggests that the elimination of stereotypical SGP's can increase the diversity of hiring outcomes without large side effects on matching efficiency, albeit at some cost for employers that were previously trying to recruit against stereotype.

2.2 Background and Conceptual Framework

2.2.1 Background

Women and men tend to work at different workplaces (see e.g., the discussion of international evidence in Card et al. 2016). While some of the reasons for this gender segregation are widely understood and accepted (such as differences in training or field of study), a long-standing concern is that employers exert discretionary preferences for one gender or the other in a particular type of job or place of work, leading to diminished opportunities, particularly for women.[5] There are many steps in the hiring process at which gender can become an obstacle. Our focus, and the focus of the information campaign we study, is at the earliest stage, when employers create and publicize information on potential job openings. In our setting these vacancies appear on a job board (see below), providing information to job seekers and intermediaries that can affect subsequent decisions about which openings are pursued.

Employers can indicate a gender preference in a newly posted vacancy in several ways. Most directly, they can indicate that they prefer male or female applicants. His-

[5] Bergmann (1974) showed how the forced segregation of women into certain occupations could lead to lower average wages for women and higher average wages for men.

torically, such preferences were widespread (see Darity and Mason, 1998 for examples in the U.S. setting). The text in a vacancy notice can also provide clues as to which type of worker an employer is looking for, particularly in German, where the use of gendered occupational titles can indicate a preferred gender (e.g. *Bauarbeiterin* for female construction worker, or a *Bauarbeiter* for male construction worker).

In the U.S., the use of gender preferences in help-wanted notices was finally outlawed by the Supreme Court in 1973, nearly a decade after the 1964 Civil Rights Act banned gender discrimination in hiring.[6] Similar changes came later in other advanced countries. Austria adopted its first Equal Treatment Act (ETA) in 1979, and outlawed gender discrimination in job advertisements in a 1985 amendment, albeit without legal sanctions. Monetary sanctions for gender preference in ads by temporary help agencies were introduced in 1992, and a June 2004 amendment extended fines to the entire private sector. Title 1 of the 2004 amendment states that recruiters *"may not advertise a position publicly or within an enterprise (company) exclusively for men or women"*, and that the job advertisements may not contain any reference from which it could be inferred that members of one sex would be favored.[7] However, exceptions could be granted in very specific cases (e.g. female care worker in shelters for women).

In the second quarter of 2005, the *Ombud for Equal Treatment* (OET), an agency that offers advice on gender equality, analyzed 36,000 newspaper job advertisements published in all major newspapers of Austria.[8] The study found that only 68% were

[6]The case, Pittsburgh Press Co. v. Pittsburgh Commission on Human Relations, 413 U.S. 376 (1973), specifically focused on the practice of listing help wanted advertisements by gender in a newspaper. See Beller (1983) for an analysis of the effect of the 1964 Civil Rights Act in the U.S. on gender segregation across occupations. There are relatively few studies of the effect of the law on gender-related outcomes: there is a much richer literature on race-related outcomes (e.g., Brown, 1984; Hersch and Shinall, 2015).

[7]This meant that advertisements had to be made gender neutral by specifying both male and female forms of any occupation, e.g. *Bauarbeiter/Bauarbeiterin.* Title I of ETA also prohibits discrimination in promotion and pay, while Title II prohibits discrimination on the basis of ethnicity, age, religion, or sexual orientation. Some provisions required changes to state laws, which had to be in effect by December 2004. Recruiters who post an SGP face a fine of 360 Euros after one warning for private sector recruiters and with no warning for public sector recruiters. Job seekers can seek damages of one month's salary if they would have received the job but for gender, and 500 Euros in case a recruiter ignored an application because of gender.

[8]The OET had conducted several previous analyses of job advertisements between 1993 and 2005, concluding that compliance with the ETA was far from universal (Lujanski-Lammer, 2006). Gender preferences may be correlated with preferences over age, ethnicity, or sexual orientation, which were also outlawed in 2004. Age requirements are recorded in our data, but this information is very coarse, only distinguishing apprenticeships from openings for trained workers. Moreover, experience requirements were not outlawed in 2004, and could substitute for age requirements. Preferences for ethnicity and

fully or partially gender neutral (Lujanski-Lammer, 2006). Over the next few months the OET conducted a major information campaign, contacting non-compliant firms and temporary help agencies, suggesting how to write gender neutral job advertisements, and notifying firms about their breach of the law. The OET also informed the help wanted sections of newspapers about the requirement of gender neutrality, so newspapers could inform employers. After the information campaign, in the fourth quarter of 2005, the OET collected a new round of data on job ads in newspapers, and found that compliance with the law had increased to about 79%, substantially higher than in Spring 2005. Below we show that by mid-2006 the vast majority of postings in the Austrian Employment Service's online job board had eliminated gender preferences.

2.2.2 Conceptual Framework

Next, we outline a simple conceptual framework, based on the model developed by Kuhn and Shen (2013) (hereafter, KS) that helps guide our empirical analysis. We assume that gender preferences are signals to job searchers (and those who assist job searchers) about a firm's preferences. Generalizing KS slightly, we assume that worker i applies to vacancy j if the subjective probability of being hired λ_{ij} exceeds some threshold τ_i.[9] We assume that λ_{ij} depends on observed information about the employer (such as the gender composition of its existing workforce), on the characteristics of the specific job, and on any stated gender preference. A stated preference can reinforce other prior information, or it can serve notice that the firm has a specific interest in a gender group that is a minority at its workplace, or in the occupation it is recruiting.

In the absence of any SGP the female share of applicants to a given job would be expected to vary with the fraction of female workers at a firm and the share of females in the relevant occupation, reflecting prior information available to searchers.[10] In the

sexual orientation are not recorded in our data. The 2005 campaign of the Ombud for Equal Treatment did not address preferences for age, ethnicity, or sexual orientation, so based on the evidence reported below regarding gender preferences, we suspect that outlawing these preferences in 2004 probably had little or no impact.

[9]KS assume that workers apply to all jobs with no SGP, and to jobs with an SGP that matches their own gender.

[10]For example, if $\tau_i \sim U[0,1]$ then the probability a specific worker applies is just λ_{ij}, so in a large population of searchers the female share of applicants would be $\mu_F \overline{\lambda}_{Fj}/(\mu_F \overline{\lambda}_{Fj} + (1-\mu_F)\overline{\lambda}_{Mj})$, where

presence of an SGP, however, the female share of applicants will depend on the degree to which searchers believe that the firm will follow its stated preference. Although we do not see application flows, we observe the gender of hired workers, and as we show below, SGP's are very strong predictors of actual hiring outcomes, over-riding other information about the firm (such as the gender composition of its existing workforce) in predicting the gender of hired workers. Similarly, KS show that the specification of a gender preference in their setting almost completely eliminates applications from the other gender.

On the demand side, a key question is why a firm would ever purposefully limit its applicant pool by specifying a gender preference.[11] KS assume that applicant screening reveals (at some cost) a candidate's match value to the firm, which has a gender-specific mean and an idiosyncratic component. The match value includes a prior assessment of the likely quality of candidates of a given gender, as well as any discriminatory taste factors. If the mean match values for men and women are far apart relative to the dispersion in the idiosyncratic component, then it is optimal for the firm to limit screening to the higher-mean gender. If the gender-specific means are relatively similar, or the idiosyncratic components are more disperse, then it is optimal to screen applicants from both genders.[12]

This model suggests that if the firm's prior assessments of match quality are correct, or it has strong tastes for one gender, then the elimination of the ability to state SGP's will increase screening costs without necessarily inducing employers to hire workers from the previously excluded gender group. If, for example, an employer has a strong distaste for women and states a preference for males when SGP's are allowed, then it seems unlikely that employer will hire women when SGP's are eliminated, even if women apply for jobs at the firm. On the other hand, if some employers are using

μ_F is the female share of all searchers and $\overline{\lambda}_{Gj}$ is the expected value of λ_{ij} for gender group G. One would expect $\overline{\lambda}_{Fj}$ to be higher for jobs at a firm with more female workers, or for more traditionally-female occupations.

[11] The same question can be asked about other qualifications specified in a vacancy, such as experience or education.

[12] Specifically KS assume that the match value of candidate i for a given opening j is a random variable $v_{ij} = v^{G(i)} + \beta\epsilon_{ij}$ where v^G is the mean for gender group G, $G(i)$ is an index function given $i's$ gender, β is a scaling factor, and ϵ_{ij} is an extreme value type 1 variate. They show that a firm's decision depends on $(v^M - v^F)/\beta$: when this term exceeds some threshold $c_U > 0$ it is efficient to screen only men, when if falls below another threshold $c_L < 0$ it is efficient to screen only women, and in the intermediate range it is efficient to screen both groups.

out-dated stereotypes to form their priors on match quality, they may be positively surprised by the quality of non-stereotypical candidates when SGP's are eliminated, and may end up filling a position with someone who would have been screened out prior to the change.

There is one case where the elimination of SGP's may lead to a *rise* rather than a fall in stereotypical hiring. That is the case of non-stereotypical SGP's. (For example, an engineering firm with only male employees may request female candidates). In such cases an SGP can override prior information that might otherwise prevent people of the non-stereotypical gender from applying. This type of SGP is similar to notices (widely used in the U.S.) that an employer is seeking a diverse application pool, but has more bite. Eliminating such SGP's may prevent job searchers or intermediaries from learning that an employer is specifically looking for a non-stereotypical candidate, and in the absence of this information they may not think it is worthwhile to apply.

2.3 Empirical Analysis

2.3.1 Data Sources

Our empirical analysis utilizes data from the job matching platform of the Austrian Employment Service, *Arbeitsmarktservice* (AMS). AMS administers Austria's income support programs (UI benefits, unemployment assistance, and related transfers) and also runs its active labor market programs. The agency has local offices in each of Austria's 104 labor market districts, at which people claiming UI have to register in person and meet regularly with staff in order to maintain benefit eligibility.

Since 1987 the AMS job-matching platform has gathered information on vacancies from firms. The platform includes almost 60% of all vacancies posted by Austrian firms, with higher coverage rates of openings in manufacturing and construction and lower rates in banking and finance (Kettemann et al., 2018; Mueller et al., 2019; Ziegler, 2021). Firms post job advertisements that list the characteristics of the open position and the desired qualifications of potential applicants. This information is entered in pre-configured fields, one of which is the preferred gender of applicants (which can be left blank).[13] Fortunately for our purposes, the preferred gender field remained in

[13]Our data include the pre-configured fields, but not the actual text posted in the online job board.

the system even after the 1984 law change because of the exemption granted by the law to certain job searches. The gender preference field allows us to examine how the 2005 information campaign, targeted to job ads, affected the use of gender preference statements in the AMS.

AMS staff use the vacancy information in their pre-selection (*Vorauswahl*) service, which routes lists of job seekers registered with AMS to firms with openings. Firms can select interviewees from among those suggested by AMS staff; they can also consider direct applications from individual workers. Vacancy information is also used by caseworkers, who can search for openings that fit the profile of a client and suggest them to the client for follow-up. If an AMS client is selected for a job opening the identity of that person is recorded in the system - a feature we use below.

The AMS job board differs from some other systems (such as the one studied by Kuhn and Shen, 2013) in the direct role played by AMS staff, though many commercial posting services use algorithms to suggest list of candidates to employers, similar to the *Vorauswahl* service. Nevertheless, as we show below, the concordance rate between SGP's in the pre-campaign period in Austria and the gender of the hired worker was similar to the concordance rate between SGP's and job applicant genders in the job board studied by Kuhn and Shen (2013). This suggests that intermediation by AMS staff and caseworkers did not necessarily diminish or magnify the role of SGP's in the matching process relative to a process driven by worker-level choices.

2.3.2 Characteristics of Vacancies

We obtained AMS vacancy data for the period September 1997 to December 2013, providing multiple years of data before and after the 2004 and 2005 events. Our dataset includes the occupation sought by the prospective employer, education requirements for the job, whether the job is full- or part-time, and whether the contract is fixed-term or open-ended. It also includes the gender preference selected by the employer and another indicator for whether the position is targeted to people aged 20+ (who have normally completed training). Finally, it includes an initial posting date and the date the vacancy was filled or closed (if the vacancy was withdrawn without being filled).

Thus, we cannot extract potential information on gender preferences from the text, contained for example in job titles.

For vacancies that were filled by AMS clients there is also an identifier of the recruited employee. We refer to such vacancies as "AMS hires".

We link vacancies with an AMS hire to the ASSD using the (anonymized) identifier of the hired employee. ASSD covers all private sector employees in Austria – about 80% of the labor force – and provides information on employment spells, including days worked at each establishment in a year and total earnings, which we convert to an average daily wage[14] (see Zweimüller et al., 2009 for more discussion of the ASSD).

In the period from September 1997 to December 2013 there were 5.2 million vacancies on the AMS system. Column 1 of Table 2.1 shows some characteristics of these vacancies, including average duration in the system (around 60 days), the share advertising for full time positions (77%) and for positions with an unlimited contract (79%), the share requiring at least upper secondary schooling (52%) and the share posted by firms with less than 4 employees (44%). Among all posted vacancies in the period up to July 2004, 27% stated a preference for males and 22% stated a preference for females (see panel B).

Unfortunately, the employer identifiers used in the AMS vacancy system cannot be linked to the ASSD. Our approach is to focus on vacancies filled by an AMS client (i.e., the AMS hires) for which we observe a person identifier that can be linked to the ASSD. For each AMS hire we search the ASSD for all employment spells involving that worker, and link the vacancy to the spell with the starting date closest to the closing date of the vacancy. We provide a full description of the matching procedure in the Appendix 2.A.

A concern for our analysis is the representativeness of vacancies filled by AMS hires. Column 2 of Table 2.1 shows the characteristics of filled vacancies, which as shown in the bottom row of the table represent about 87% of posted vacancies. Relative to the characteristics in column 1 there are few major differences, suggesting that *filled* vacancies are similar to the broader population of filled or withdrawn vacancies. Next, in column 3, we present the characteristics of the 1.7 million filled vacancies with an AMS hire. This subset, which represent 26.6% of all filled vacancies, have a slightly lower mean duration (44 days vs. 58 days), are slightly more likely to advertise for a position with an unlimited contract, and are slightly less likely to require high-school level education (46% vs. 52%). In the pre-campaign period AMS-filled vacancies were

[14]The ASSD database does not contain information on working hours, though we know the part-time status of a vacancy.

Table 2.1: Vacancy Characteristics

	Subsample			
	All Outflows	**Hired**	**Hired AMS**	**Matched**
	(1)	(2)	(3)	(4)
Panel A: Vacancy Characteristics				
Avg Vacancy Duration (days)	52.72	51.86	40.97	40.91
	(30.23)	(29.89)	(28.45)	(28.37)
Full Time Position (share)	0.77	0.77	0.78	0.78
	(0.42)	(0.42)	(0.41)	(0.41)
Unlimited Contract (share)	0.79	0.78	0.81	0.82
	(0.41)	(0.41)	(0.39)	(0.39)
Small Firms (share)	0.44	0.44	0.44	0.43
	(0.50)	(0.50)	(0.50)	(0.50)
Upper Secondary Education (share)	0.52	0.52	0.46	0.47
	(0.50)	(0.50)	(0.50)	(0.50)
Panel B: SGP - pre-campaign period				
Preference for Men (share)	0.26	0.26	0.30	0.30
	(0.44)	(0.44)	(0.46)	(0.46)
Preference for Women (share)	0.22	0.22	0.23	0.23
	(0.41)	(0.42)	(0.42)	(0.42)
Observations	4,998,146	4,333,444	1,151,996	987,271
Observations - pre-campaign period	1,622,911	1,419,630	462,450	397,496

Notes: This table reports means and shares of selected characteristics of vacancies in four sub-samples. Period 1997-2013. Pre-campaign period refers to the period 1997-2004. The sub-sample "all" refers to all vacancies, "Hired" is for vacancies that are filled, "Hired AMS" is the sub-set of vacancies filled through AMS, and "Matched" is the sub-set of the AMS hires (in column 3) that we matched with a firm in the Austrian Social Security Database (ASSD). Vacancy Duration refers to "*Vacancy Closing Month (last day) - Vacancy Posting Date*" as the employment spell starting date is not available for non matched vacancies.

Source: AMS-ASSD data, own calculations.

also somewhat more likely to state a preference for a male (31% versus 26% of filled vacancies), whereas the fraction stating a preference for a female is about the same as in the overall population of vacancies (23% versus 22%).

More information is presented in Appendix Tables 2.A2 and 2.A3, where we show 1-digit industry and occupation shares for the same sets of vacancies summarized in Table 2.1. Jobs in mining, manufacturing, construction, public administration and health are somewhat over-represented in the set of vacancies with an AMS hire, while jobs in hotels and restaurants, finance, real estate and business services, and health care are under-represented. At the occupation level, positions for clerks, operatives, and unskilled (elementary) occupations are over-represented among vacancies filled by AMS hires, while those for managers, professional and technical workers, and shop and sales workers are under-represented.

Finally, column 4 presents the characteristics of the subset of vacancies with an AMS hire that we successfully matched to an employer in the ASSD data. We lose about 14% of potential matches because of missing or incomplete data, or because there is no new employment spell in the ASSD that is an obvious match for the filled vacancy in the AMS data. Reassuringly, the characteristics in columns 3 and 4 are quite similar (as are the corresponding industry and occupation shares shown in Appendix tables 2.A2 and 2.A3), suggesting that the matching process is approximately random.

While AMS-filled vacancies do not have exactly the same characteristics as the more general population of filled vacancies, we believe the differences are modest. In the remainder of the paper we therefore use the sample described in column 4 of Table 2.1 as our main analysis sample. As a robustness check we use the characteristics of vacancies in column 1 to construct weights that represent the inverse probability that a vacancy was filled by an AMS client, and re-estimate our main models using these weights (see below). Reassuringly we find relatively small discrepancies from our unweighted models.

2.3.3 Use of Stated Gender Preferences Prior to 2005

We begin our analysis by looking at what happened to the fraction of vacancies with a stated preference after the introduction of fines for noncompliance with the Equal Treatment Act in June 2004, and the information campaign in early 2005, Figure 2.1

(Recall that the AMS system allowed employers to select a preferred gender even after they were declared illegal to accommodate exemptions from the law).

Figure 2.1: SGP in time

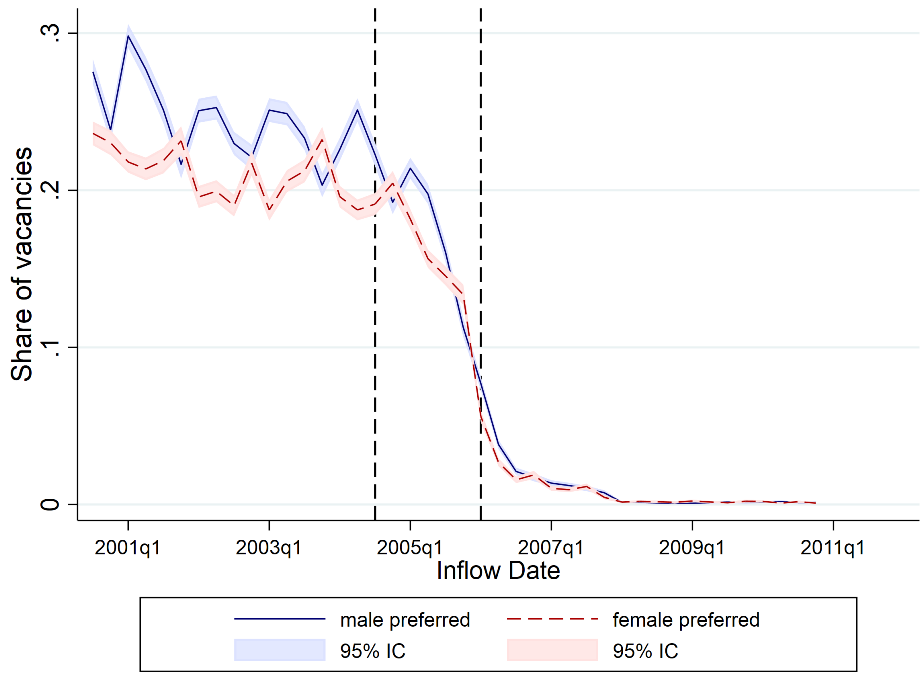

Notes: This figure shows that share of posted vacancies in the AMS job board system that specify a preference for females (dashed line), or males (solid line). The vertical dashed lines depict the date of a law change introducing sanctions for posting gender preferences (July 2004), and the timing of an information campaign to alert employers and newspapers about the law (Spring 2005). Since the law allows some exemptions, the system continued to allow employers to post preferences throughout the sample period.

Source: AMS-ASSD data, own calculations.

In the years prior to 2004 the fractions of AMS-filled vacancies with a stated preference were fairly stable, apart from an obvious seasonal pattern. At the July 1 2004 effective date of the amendment introducing fines there is no obvious trend break in either rate, suggesting that the introduction of sanctions had little direct effect. In the second quarter 2005, however, the shares of vacancies with male or female preferences both begin to decline sharply, and by mid-2006 only a small share of vacancies express a gender preference. The trend break in 2005 coincides exactly with the information campaign conducted by the OET from the second to the fourth quarter of 2005. The coincidence of timing suggests that the OET's campaign effort was responsible for the near elimination of SGP's, though this effect has to be understood as coming after sanctions were introduced into the Equal Treatment Act, and in light of evolving attitudes toward gender.

Given the patterns in Figure 2.1, we consider the period up to and including 2004 as the "pre" period for our event studies and difference-in-differences models, and the period from 2006 onward as the "post" period. For our main models we drop 2005 (the year of transition), although in our event study graphs we retain these observations.

How did AMS-filled vacancies stating a male or female gender preference differ from those with no SGP? We investigate these differences in Table 2.2, which uses data from the 2000-2004 (pre-campaign) period. About 21% of vacancies stated a gender preference for females (column 1), 26% for males (column 3), and 53% stated no preference (column 2). During these years females made up about 46% of all people hired for AMS-filled vacancies (see column 4 of the table). The share of women hired for jobs with no SGP was quite similar to their overall share, but for vacancies with a female SGP, 96.8% were filled by a female, while for vacancies with a male SGP, only 2.6% were filled by a female. Thus, compliance with SGP's was extremely high.

Table 2.2: Descriptive Statistics on New Jobs Associated with Filled Vacancies

	SGP			
	Female SGP (1)	**No SGP** (2)	**Male SGP** (3)	**Total** (4)
Share of vacancies (share)	0.213	0.529	0.259	
Panel A: Outcomes				
Share of women hired (share)	0.967	0.462	0.028	0.457
Vacancy filling time (mean)	30.70	34.53	30.65	32.71
Vacancy filling time (median)	20	23	19	21
Log wage of the hire (mean)	3.58	3.83	3.98	3.81
Job duration (mean)	425	397	338	388
Job duration (median)	195	177	144	171
Panel B: Context				
Share of women in the firm (share)	0.691	0.480	0.247	0.464
Share of women in occupation (share)	0.659	0.496	0.271	0.473

Notes: This table reports means, medians and shares of selected characteristics of vacancies in the matched sample by Stated Gender Preference (SGP) during the 2000-2004 (pre-campaign) period. Vacancy filling time refers to the difference in days between the starting date of the new job associated with the vacancy and the posting date of the vacancy.

Source: AMS-ASSD data, own calculations.

We also report the mean and median duration of filled vacancies, the mean log daily wage associated with the newly created jobs, and the duration of these jobs (mean and median). We can see a couple of interesting patterns that will carry through our entire analysis. First, job openings with an SGP *for either gender* tend to be filled a little faster.[15] Second, jobs created by filling a vacancy with a female SGP – which are nearly all held by women – pay relatively low wages but have longer durations than than those created by filling vacancies with no SGP. (Some of this pay gap could be due to differences in hours of work). Symmetrically, jobs created from a vacancy with a male SGP – which are nearly all held by men – pay relatively high wages and have a shorter duration.

Figure 2.2: Gender Mix Within Firms

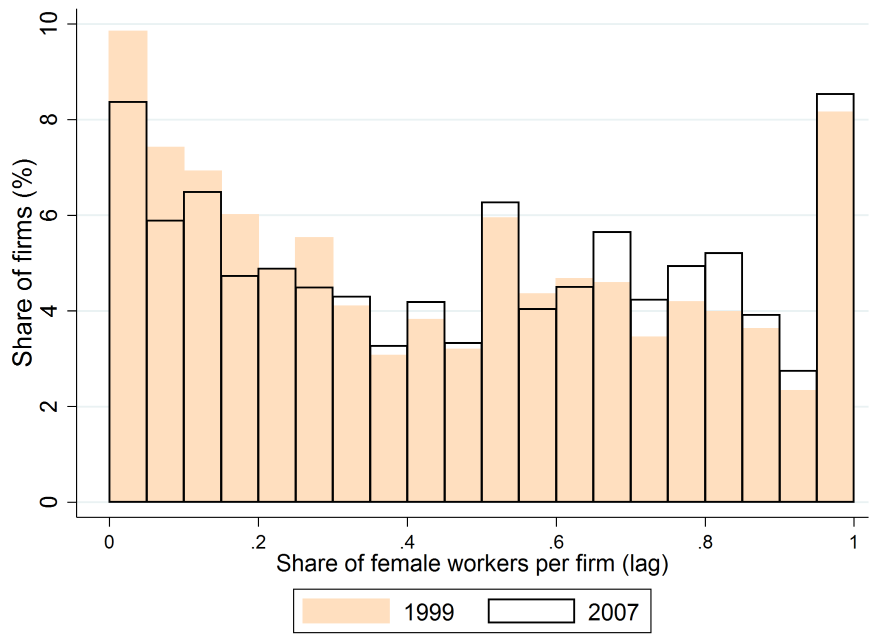

Notes: This figure displays the histogram of the share of women in firms for 1999, and 2007. Histograms are estimated using .05 bin size.

Source: ASSD

The key question motivating our work is the extent to which SGP's tend to perpetuate gender segregation across workplaces. Figure 2.2 shows histograms of the share of women in each establishment in our sample in 1999 (5 years before the campaign) and

[15]Vacancy filling times reported in Table2.2 are based on the elapsed time from the day the vacancy was posted and the day the new employee started work. The times reported in Table 1, which pertain to all filled vacancies including those we cannot match to the ASSD, are based on the number of days between the posting date of the vacancy and the last day of the month in which it was removed from the AMS job board.

2007 (3 years after). In both years the distribution of the share of women is inverse-U-shaped, with a relatively large share having <10% female workers and another relatively large group having 100% females.[16] Comparing histograms for the two years we can see clear evidence that the degree of gender segregation across firms in Austria fell somewhat between 1999 and 2007. The average share of women also increased over this period, however, making it hard to credit the change solely to the elimination of SGP's.

How did the use of stated gender preferences in the years prior to the campaign vary with the female share of employees? To provide some initial descriptive evidence, we construct a measure of the fraction of female employees at the firm that posted the vacancy in the year prior to the posting date. We then use data on vacancies in the pre-campaign period to estimate the average fraction of vacancies with either no SGP, or a male or female SGP, among workplaces in different ranges of the lagged female share. The results, presented in Figure 2.3, show the expected pattern.

Figure 2.3: Share of Job Ads by Gender Preference

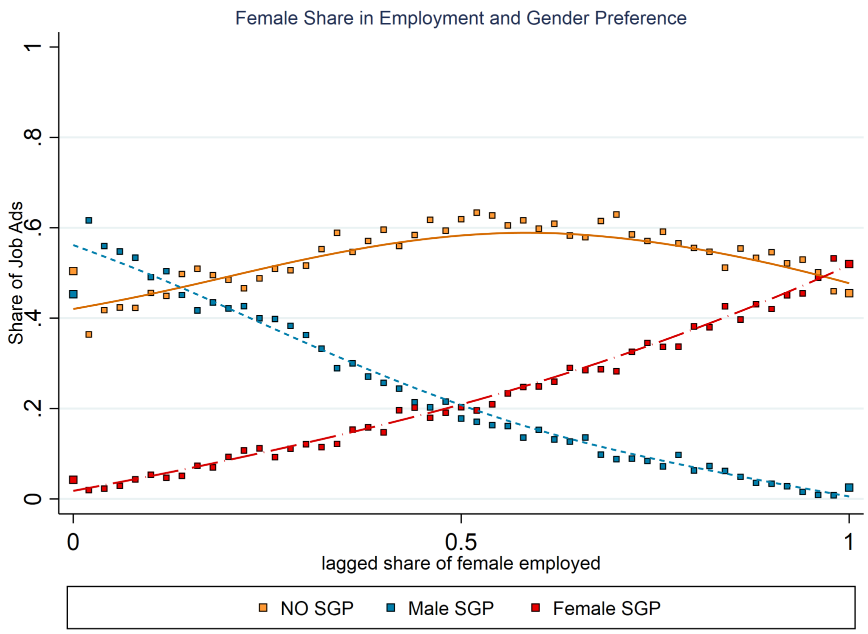

Notes: This Figure shows the share of vacancies that specify preferences for females (red line), males (blue line), or no gender preference (yellow line), by gender composition of the firm in the year prior to the job advertisement. Data are drawn from 2000-2004.
Source: AMS-ASSD data, own calculations.

[16] The "spikes" in the distributions in Figure 2.2 are due in part to firms of size 1 or 2. See D'Haultfœuille and Rathelot (2017) for a discussion of small unit bias in measuring characteristics of firms.

Workplaces with few women (toward the left of the figure) tended to post vacancies with either a *male* SGP (about 60% of the time) or *no* SGP (around 40% of the time). Symmetrically, workplaces with few men (toward the right of the figure) tended to either post vacancies with a *female* SGP (about 60% of the time) or *no* SGP (around 40% of the time). Thus, when they were allowed, SGP's tended to be used to reinforce existing gender segregation patterns. There was also some limited use of SGP's to recruit women at firms with <50% female share, and to recruit men at firms with <50% males. In our analysis below we refer to such preferences as non-stereotypical SGP's.

Next we look at the relationship of SGP's to ***hiring outcomes***: Figure 2.4 shows the average share of women actually hired by firms with different (lagged) female employment shares, conditional on the type of SGP (male, female, or none).

Figure 2.4: Share of Female Hired by Gender Preference

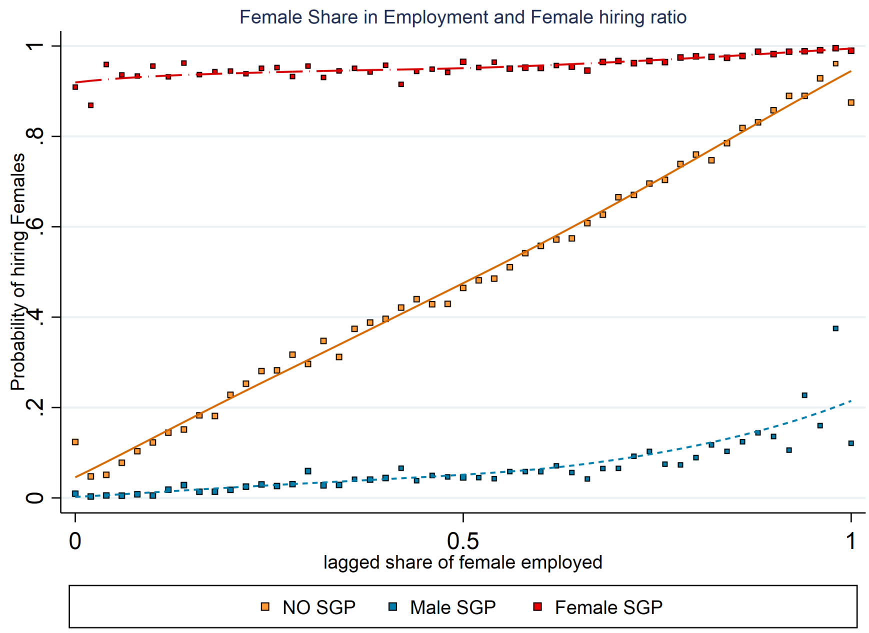

Note: This Figure shows shows the share of females hired by gender composition of the firm in the year prior to the hire, and by whether the vacancy stated a preference for females (red line), males (blue line), or neither. The data are from 2000-2004.
Source: AMS-ASSD data, own calculations.

The figure shows stark differences in hiring outcomes for these three types of vacancies, even controlling for workplace composition. In the absence of any SGP (the yellow line in the figure), the fraction of females hired is approximately linear in the lagged female share, with an intercept that is just slightly above 0 and slope that is just slightly flatter than 1. For vacancies with a **male** SGP (blue line), however, the fraction of females hired rises only slightly with the workplace share of females, reaching about 20% even at firms that had all female workers in the previous year. In contrast, for vacancies with a **female** SGP (red line), the fraction of females hired starts at around 90% at workplaces with no women in the previous year, and rises slightly to essentially 100% at all-female workplaces.

The remarkable "Z" shape of Figure 2.4 suggests that, when SGP's were allowed, they exerted a strong impact on hiring outcomes. Workers hired to fill vacancies with no SGP tended to perpetuate the status quo gender composition, with only a small probability of hiring a female at all-male firms (around 5%), or of hiring a male at all-female firms (also around 5%). But nearly all the new hires for vacancies with a female SGP were female (even at all-male workplaces) and nearly all the new hires for vacancies with a male SGP were male (even at all-female workplaces).

A potentially confounding factor in the interpretation of this Z pattern is the role of occupation. Many occupations are highly specific to one gender (Cortes and Pan, 2017), so the female share of a firm's workforce is partly driven by occupation structure. We would expect new hires to obey the same structure, leading to a strong positive relationship between the lagged share of females and the average share of females in new hires. If employers tend to use a male SGP when they are recruiting for mainly male occupations and a female SGP for mainly female occupations, however, we would see a close match between SGP's and the gender of hired workers, even if SGP's have no effect on application or hiring decisions.

To isolate the effects of SGP's while controlling for occupation, we used data from the EU Standard of Income and Living Conditions (SILC) data base, **excluding Austria**, to assign the share of females in each of 100 occupation groups (International Standard Classification of Occupations, ISCO, 2-digit major sub-groups). We then linked these gender shares to the vacancies in our data set, and divided vacancies in the pre-campaign period into 3 groups: those for mainly male occupations (with less than 30% female workers on average); those for mainly female occupations (with more than 70% female on average); and those for mixed-gender (or "gender neutral") occupations. Finally, we constructed separate versions of Figure 2.4 for each of the three groups of occupations.

As shown in Figure 2.5, all three types of vacancies exhibit a Z pattern, though the

Figure 2.5: Role of Female Share in Occupation

(a) Female Occupations

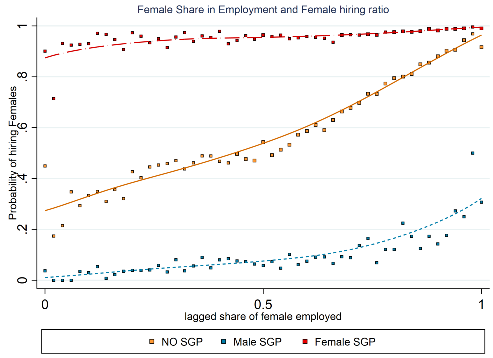

(b) Gender Neutral Occupations

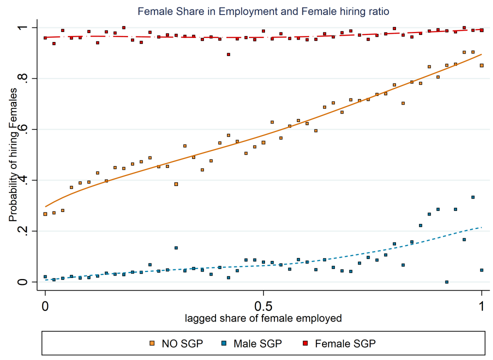

(c) Male Occupations

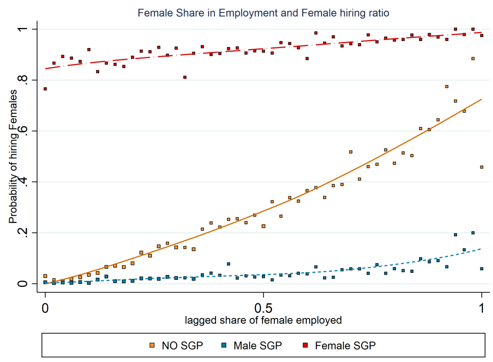

Note: This Figure shows shows the share of females hired by gender composition of the firm in the year prior to the hire, and by whether the vacancy stated a preference for females (red line), males (blue line), or neither (yellow line), for three groups of vacancies: those advertising for mainly female occupations (panel a), mainly male occupations (panel c) or occupations with a mixed gender composition (panel b), The data are from 2000-2004.

Source: AMS-ASSD data, own calculations.

precise shape of the Z – particularly the diagonal element representing vacancies with no SGP – varies by occupation group. In the absence of an SGP, the average fraction of women hired to fill *female* occupations starts at around 25% at all-male workplaces and rises to 100% at all-female workplaces. Symmetrically, the average fraction hired to fill *male* occupations starts at 0 at all-male workplaces and rises to around 75% at all-female workplaces. The average fraction hired to fill *gender neutral* occupations lies between these extremes, starting at around 25% and increasing to around 85%. Thus, for vacancies with no SGP, controlling for occupation flattens the relationship between the lagged female share and the probability of hiring a female. For vacancies with male or female SGP's, however, most of the new hires match the stated gender preference, regardless of whether we control for the gender mix of the occupation or not, and regardless of the gender mix of the workplace.

Stated gender preferences appear to influence hiring outcomes, but is there any evidence that they affect hiring efficiency (as is assumed in the KS model)? To assess this we look at vacancy filling times. Columns 1-4 of Table 2.3 present a series of simple models that relate the days required to fill a vacancy to the use of SGP's. To help provide a context for these models we also present a parallel set of models for the event that a female was hired to fill the vacancy (columns 5-8). We look at all vacancies (columns 1 and 5), and separately at vacancies for mainly male occupations (columns 2 and 6), for occupations with a mix of male and female workers (columns 3 and 7), and for mainly female occupations (columns 4 and 8).

We classify workplaces as "female" (more than 50% females in the previous year, with label "F workplace") or "male" (more than 50% males in the previous year, , with label "M workplace") and include as the main variables of interest indicators for a male or female SGP, interacted with dummies for mainly male or mainly female workplaces. All the models include fixed effects for occupation and industry, time effects, and dummies for 51 different intervals of the lagged share of females at the workplace.

Looking at the vacancy filling time results we see two broad patterns. First, the use of SGP's tends to (if anything) reduce vacancy filling times. Second, SGP's for the gender that is the opposite of the workplace majority tend to have the largest negative effects. The large negative effects of a female SGP in filling an opening for mainly male occupation at a mainly male workplace (-2.406 days, in the 4th row of column 2) and of a male SGP in recruiting for a mainly female occupation at a mainly female workplace

(-6.517 days, in the first row of column 4) are particularly interesting. In both cases, the employer is stating a preference for a non-stereotypical gender (different than the majority of the occupation and of the existing workforce), and yet the time to fill the vacancy is reduced. As shown in the corresponding models for the gender of the newly hired worker (columns 6 and 8), such preferences have large effects on the probability of hiring the preferred gender, consistent with the patterns in Figure 2.5. Thus, it appears that employers with non-stereotypical preferences could easily find workers to match those preferences. Taken together, the results on filling times and gender outcomes suggest that in the pre-campaign period there were more females looking for jobs in mainly male occupations, and more males looking for jobs in mainly female occupations, than were demanded by the relatively small share of employers who were seeking to recruit them.

Table 2.3: Vacancy Filling Time and Female Hiring

	OLS Estimation							
Dependent Variable:	*Vacancy Filling Time - in days*				*Female hire*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
M SGP in F workplace	-2.140***	-2.960***	2.038	-5.684***	-0.378***	-0.306***	-0.348***	-0.443***
[Non-stereotypical Male PGP]	(0.446)	(0.836)	(1.504)	(0.752)	(0.005)	(0.010)	(0.017)	(0.009)
M SGP in M workplace	-1.212***	-0.394	-0.739	-4.637***	-0.139***	-0.081***	-0.231***	-0.277***
[Stereotypical Male PGP]	(0.289)	(0.375)	(0.831)	(1.019)	(0.002)	(0.003)	(0.009)	(0.011)
F SGP in F workplace	-3.042***	-2.662***	-2.099**	-3.148***	0.240***	0.447***	0.182***	0.227***
[Stereotypical Female PGP]	(0.296)	(1.022)	(1.053)	(0.356)	(0.003)	(0.012)	(0.011)	(0.004)
F SGP in M workplace	-3.274***	-4.440***	-4.572***	-2.196***	0.469***	0.673***	0.335***	0.423***
[Non-stereotypical Female PGP]	(0.405)	(0.807)	(0.931)	(0.770)	(0.004)	(0.009)	(0.011)	(0.008)
All Vacancies	✓				✓			
Male Occupations		✓				✓		
Neutral Occupations			✓				✓	
Female Occupations				✓				✓
Observations	255,969	102,630	48,046	105,293	255,969	102,630	48,046	105,293

Notes: OLS estimation of the effect of the SGP on vacancy filling time and hiring gender. Period 2000-2004 (pre-campaign). Controls include occupation, industry and firm gender composition fixed effects as well as year FE. Beta coefficients reported and robust standard errors in parentheses. Columns (1) and (5) concerns all observations. In column (2) and (6) only job openings in female occupations are used, in column (3) and (7) in neutral occupations and in column (4) and (8) in male occupations only. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

To summarize: in the years prior to the 2005 campaign, nearly one-half of vacancies were posted with a stated gender preference. Most employers were stating stereotypical gender preferences that tended to reinforce gender segregation, though a minority were using non-stereotypical preferences to recruit against type. Moreover, firms that stated a gender preference were very likely to recruit a worker of the preferred gender, and were able to fill the vacancy somewhat faster, particularly in cases where the stated preference was for a gender different from the majority of the occupation and the firm's existing workforce. What we cannot tell from the observational data is the extent to which firms that advertised for a particular gender were willing to hire people of the opposite gender, and if so whether the resulting matches would be much different from the ones initiated prior to the campaign. These are the questions to which we now turn.

2.4 Effects of the Campaign on Stated Gender Preferences

In the post-campaign years we cannot see the gender preferences that employers would have stated in the absence of the campaign.[17] Instead, we develop a prediction model that identifies vacancies that would have been likely to have a male SGP, a female SGP, or no SGP in the pre-campaign period. We refer to these as "predicted gender preferences", or PGP's. Then we use this model to classify vacancies *before and after the campaign* into the three groups, and examine changes in outcomes for the three sets of vacancies after the campaign.[18]

2.4.1 Classifying Vacancies by Predicted Gender Preference

We use data on AMS-filled vacancies in the 2000-2004 period to estimate prediction models for the use of gender preferences. Specifically, we use information on the industry I_j (in 86 categories) and occupation O_j (in 100 categories) associated with the j^{th}vacancy, together with data on the lagged female share at the posting establishment F_j (in 51 intervals). In our first approach, we define a variable S_j for each vacancy equal

[17]Kuhn and Shen (2021) are able to compare applications to vacancies that were initially posted before gender preferences were removed from the job board they study, and remained posted after their removal. This design holds constant all features of the vacancy except the SGP's.

[18]This approach is related to designs that classify subgroups based on the likelihood of being affected by a treatment, as in the "gap design" for studying minimum wages (e.g., Card, 1992). Botosaru and Gutierrez (2018) discuss identification and estimation in related studies using repeated cross sectional data in which treatment status is only observed after treatment.

to 1 if it has a female SGP, equal to -1 if it has a male SGP, and equal to 0 otherwise. (Note that $S_j = S_j^f - S_j^m$, where S_j^f is a dummy for a stated female preference and S_j^m is a dummy for a stated male preference). We then calculate the leave-out mean (LOM) value $\overline{S}_{\sim j}$ for vacancies in each (I, O, F) cell and assign this back to vacancies in the pre-campaign and post-campaign periods.[19] Finally, we classify a vacancy as likely to have a female SGP (indicated by a dummy D_j^f) if $\overline{S}_{\sim j}$ is above some threshold:

$$D_j^f = 1[\overline{S}_{\sim j} > c_1]$$

and likely to have a male SGP (identified by D_j^m) if $\overline{S}_{\sim j}$ is below some other threshold:

$$D_j^m = 1[\overline{S}_{\sim j} < c_2].$$

Vacancies with neither condition true are classified as likely to have no SGP. We select the cutoffs c_1 and c_2 so that the predicted shares of vacancies with a female or male SGP match the actual shares with these preferences in the pre-campaign period. We refer to the dummies D_j^f and D_j^m as ***predicted gender preferences*** (PGP's). Note that we cannot assign leave out means to (I, O, F) cells with only 1 vacancy. Such vacancies are omitted from the sample used in all subsequent analysis.

This classification procedure is equivalent to calculating the leave-out mean fractions of vacancies in each (I, O, F) cell with female and male SGP's, $\overline{S}^f_{\sim j}$ and $\overline{S}^m_{\sim j}$, respectively, then classifying a vacancy based on whether $\overline{S}^f_{\sim j} - \overline{S}^m_{\sim j}$ is above some threshold (for a female SGP) or below some other threshold (for a male SGP). Thus, holding constant $\overline{S}^f_{\sim j}$, vacancies in a cell with a higher value of $\overline{S}^m_{\sim j}$ are less likely to be classified as a female PGP. Likewise, holding constant $\overline{S}^m_{\sim j}$, vacancies in a cell with a higher value of $\overline{S}^f_{\sim j}$ are less likely to be classified as a male PGP. We believe this ordered structure makes sense: an employer may be close to indifferent between posting a job with a particular gender preference or no gender preference, but would hardly ever be indifferent between posting a job with a male SGP or a female SGP.

In our second approach, we estimate a Random Forest (RF) model using 4-fold cross validation to classify SGP's using the same categorical I, O, F data for the 2000 to 2004

[19]Since the post-campaign vacancies are never used in calculating the SGP shares, the LOM assigned to these vacancies includes all the data.

period.[20] We construct 100 trees and take the majority vote for the classification of each vacancy as our final classification.

We show the prediction success rates for our two different approaches in Appendix Figure 2.C1. We define the success rate in a given year as the fraction of vacancies in that year with a given SGP type that were predicted to have that SGP type (i.e., $E[D_j^f|S_j^f = 1]$ and $E[D_j^m|S_j^m = 1]$). The data for 1999-2004 are "in-sample" predictions, though by using a leave-out mean (in our first approach) or k-fold cross validation (in our second) we reduce the risk of in-sample over-fitting. The data for 2005 and 2006 are "out-of-sample" predictions, since data from these years are not used in our models. We stop after 2006 because by 2007 most employers were no longer using SGP's. For comparative purposes we also show predictions that use the entire sample in either approach in the two right hand panels of the figure.

The upper left panel of the figure shows that the fraction of SGP's correctly predicted by our first approach in the in-sample period is just over 60% for vacancies with male SGP's and no gender preference, and around 55% for those with female SGP's. All three rates evolve fairly smoothly between the in-sample and out-of-sample period. The lower left panel shows that the RF classifier has somewhat higher success rates for male SGP's (around 65%) and no SGP (close to 70%), but a lower success rate for female SGP's (around 41%). In the out-of-sample period the RF's prediction success rate also falls more abruptly, suggesting some tendency for over-fitting despite the use of cross validation in building the classifier.

The patterns in two right hand panels illustrate the importance of not using data from a given vacancy in forming a prediction of its SGP. For both our classifiers, using the own-vacancy data leads to higher prediction success in the in-sample period but a noticeable degradation in performance in the out-of-sample period, indicative of over-fitting bias.

2.4.2 Measuring Effects on Hiring Outcomes

We evaluate the effects of the OET campaign on the use of stated gender preferences using a simple difference in differences approach that compares post-campaign out-

[20]Note that this approach can in principal calculate predicted preferences even for (I, O, F) cells with only 1 vacancy. However, for comparison with the LOM approach we drop these from the sample.

comes to pre-campaign outcomes for three groups of vacancies: those with a predicted female preference ($D_j^f = 1$), those with a predicted male preference ($D_j^m = 1$) and those with no predicted preference ($D_j^f = D_j^m = 0$). Let y_j represent some outcome associated with the vacancy (e.g., the gender of the hired candidate or the wage paid to the new worker). Then we fit models of the form:

$$y_j = \beta_0 + \beta_1 D_j^f + \beta_2 D_j^m + \lambda_1 D_j^f Post_j + \lambda_2 D_j^m Post_j + X_j \delta + \varepsilon_j \tag{2.3.1}$$

where $Post_j$ is an indicator for vacancy j being listed in 2006 or later, and X_j is a set of control variables, including time effects (which absorb the main effect of $Post_j$), and dummies for industry, occupation and the female share of employees at the posting employer in the previous year. Recognizing that D_j^f and D_j^m are based on a first-stage prediction model, we bootstrap the standard errors for the estimated coefficients from this model.

The coefficients λ_1 and λ_2 in equation (1) measure the changes in the outcome y between the pre-campaign and post-campaign period for vacancies that were predicted to have a female or male SGP, respectively, *relative to the trend for vacancies that were predicted to have no expressed gender preference*. Since the law did not directly effect recruiters that would have posted vacancies with no gender preferences, the outcomes for their vacancies are a natural comparison group. We note, however, that general equilibrium effects could potentially change the types of workers who apply for these vacancies.

Since we are using predicted rather than actual gender preferences, equation (1) can be interpreted as the reduced-form from a two-stage-least-squares procedure in which we estimate the first stage models using only data from the pre-campaign period, and allow the effects of the endogenous variables in the outcome equation to vary between the pre-campaign and post-campaign periods. To formalize this reasoning, assume that the true model generating outcome y is:

$$y_j = \alpha_0 + \alpha_1 S_j^f + \alpha_2 S_j^m + \theta_1 S_j^f Post_j + \theta_2 S_j^m Post_j + X_j \gamma + \varepsilon_j \tag{2.3.2}$$

where S_j^f and S_j^m indicate actual SGP's in the pre-campaign period and *desired* SGP's in the post-campaign period (i.e., the preferences that employers would have stated if there was no campaign). In this model, θ_1 and θ_2 measure the changes in outcome y between vacancies in the pre-campaign period that were posted with given SGP and vacancies in the post-campaign period that would have been posted with that gender preference if such preferences were allowed. Notice that these effects are conceptually the same as the effects measured in the design of Kuhn and Shen (2021), which compares applications for the same vacancy after SGP's are eliminated.

Assume that the actual or desired SGP's are related to PGP's by a pair of simple models with constant coefficients between the pre- and post-campaign periods:

$$S_j^f = \pi_0 + \pi_1 D_j^f + \pi_2 D_j^m + X_j \pi_x + \xi_j^f \tag{2.3.3}$$

$$S_j^m = \psi_0 + \psi_1 D_j^f + \psi_2 D_j^m + X_j \psi_x + \xi_j^m \tag{2.3.4}$$

where ξ_j^f, ξ_j^m are prediction errors. Here π_1 represents the increment in the probability of an actual female SGP if the vacancy has a predicted female SGP relative to the omitted category of no predicted SGP, and π_2 represents the increment in the probability of an actual female SGP if the vacancy has a predicted male SGP, again relative to the case where it is predicted to have no SGP. Thus we expect π_1 to be positive and π_2 to be negative (though small in magnitude). Similar reasoning suggests that ψ_2 will be positive and $\psi_1$will be negative.

Combining equation (2) with (3) and (4) shows that the difference-of-differences coefficients in (1) are:

$$\lambda_1 = \theta_1 \pi_1 + \theta_2 \psi_1 \tag{2.3.5}$$

$$\lambda_2 = \theta_1 \pi_2 + \theta_2 \psi_2. \tag{2.3.6}$$

Notice that if we ignore π_2 and ψ_1,then λ_1 is an attenuated version of θ_1 and λ_2 is an attenuated version of ψ_2, where the attenuation factors reflect the fractions of predicted vacancies with female and male preferences that actually have these preferences (conditional on the $X's$). More generally we expect π_2 and ψ_1 to be small negative numbers, so this intuition remains roughly correct.

Columns 1 and 3 of Table 2.4 present estimates of equations (3) and (4) using the observed SGP's in the 1999-2004 period and predictions from our first (leave-out-mean based) classification model. We see that $\pi_1 = 0.173$ and $\pi_2 = -0.035$, while $\psi_1 = 0.192$

and $\psi_2 = -0.027$. Thus, controlling for industry, occupation, and the firm's lagged gender composition, having a PGP of one gender raises the probability of an SGP of that gender by 17-19 percentage points relative to a vacancy with neither PGP, while having a PGP of the opposite gender lowers the probability by about 3 percentage points relative to the comparison group of neither PGP. (Columns 2 and 4 present estimates for vacancies by workplace context. We come back to these models when we discuss results that compare stereotypical and non-stereotypical preferences in Section 2.4.2.)

Table 2.4: Relationship between Predicted and Actual Gender Preferences in Pre-campaign Period

	OLS Estimation			
Dependent Variable:	*Male SGP*		*Female SGP*	
	(1)	(2)	(3)	(4)
No PGP	omitted	omitted	omitted	omitted
Male PGP	0.192***		-0.035***	
	(0.004)		(0.001)	
Female PGP	-0.027***		0.173***	
	(0.002)		(0.004)	
M PGP in F workplace		0.266***		-0.077***
[Non-stereotypical Male PGP]		(0.009)		(0.005)
M PGP in M workplace		0.181***		-0.029***
[Stereotypical Male PGP]		(0.004)		(0.001)
F PGP in F workplace		-0.012***		0.164***
[Stereotypical Female PGP]		(0.002)		(0.004)
F PGP in M workplace		-0.068***		0.197***
[Non-stereotypical Female PGP]		(0.004)		(0.007)
Observations	217,568	217,568	217,568	217,568

Notes: OLS estimation. Regression of the actual SGP's on the predicted gender preferences (PGP's) using data from 2000-2004 only. Controls include occupation, industry and firm gender composition fixed effects as well as year FE. Beta coefficients reported and robust standard errors in parentheses. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

Source: AMS-ASSD data, own calculations.

These estimates imply that the coefficients λ_1 and λ_2 are substantially attenuated – by a factor of roughly 80% – relative to the effects we would estimate if we could see the desired stated preferences of job recruiters in the post-campaign period. In other words, the true effects of eliminating stated gender preferences may be 5 times larger than the estimates we obtain from equation (1).

Before proceeding we note two other methodological points. First, we are assuming that the intervention of interest is the elimination of firms' abilities to declare their gender preferences in job advertisements. One concern with this assumption, particularly for our analysis of wage outcomes, is that the 2004 law change or the OET campaign may have also affected firms' pay decisions. While the OET's campaign did not directly address pay disparities, it is possible that it heightened awareness of gender-related pay gaps. Any differential changes in pay policies between firms that were more or less likely to use gender preference will be confounded in our difference in differences analysis.

Second, the framework of equations 2-4 makes clear that we are identifying the effect of stated gender preferences using interactions between occupation, industry, and the gender composition of the firm's workforce as instruments (since main effects for these 3 variables are included in all our models). Given the stark patterns in Figures 3 and 4 we believe this is a plausible strategy. As a check, we have re-estimated our main models for the gender of the newly hired worker adding industry×year effects, occupation×year effects, or workplace gender composition×year effects, and find that the implied impacts are quite similar (Figure 2.D4 in Appendix).

Effects on the of Probability of Hiring a Female Candidate

We begin by using this simple framework to estimate the effects of eliminating the use of stated gender preferences on the probability of hiring a female worker. Before reporting difference-in-difference results, we present event study graphs in Figure 2.6. These are based on a slightly more general specification:

$$y_j = \beta_0 + \beta_{1t} D_j^f + \beta_{2t} D_j^m + X_j \delta + \varepsilon_j \tag{2.4.1}$$

with separate coefficients on D_j^f and D_j^m in each year. We estimate the pair of coefficients (β_{1t}, β_{2t}) for each year of our sample, then normalized the estimates relative to 2004 by forming $\widehat{\beta}_{1t} - \widehat{\beta}_{1,2004}$ and $\widehat{\beta}_{2t} - \widehat{\beta}_{2,2004}$. The left panel of the figure plots the values of the normalized series using our LOM classifier to predict SGP's, while the right panel shows estimates based on the Random Forest classifier.

The results using either classifier suggest that in the period up to 2004 the fraction of females hired for vacancies with a predicted female or male gender preference (PGP) were relatively constant. Starting in 2005, however, there was a clear shift, with a rise in the fraction of women hired to fill vacancies that were predicted to have a male PGP (blue lines in each panel) and a fall in the fraction of women hired to fill vacancies that were predicted to have a female PGP (red lines in each panel). The magnitudes of the shifts differ somewhat between the two prediction models, with the LOM classifier suggesting a roughly 2-3 percentage point rise in the fraction of women hired for openings with a predicted male PGP, and a 1-2 points fall in the fraction hired for openings with a predicted female PGP, and the RF classifier suggesting a smaller rise in the former group (on order of 1-2 percentage points) but a larger fall for the latter group (2-3 percentage points). Note that the <u>difference</u> in the probability of hiring a female worker between vacancies with predicted male and female preferences (i.e., the difference between the blue line and the red line) increased by 3-5 percentage points after the OET campaign using either classifier.

Figure 2.6: Stated Gender Preference Affects Female Hiring: Event History Results

A. Leave-Out-Mean B. Random Forest

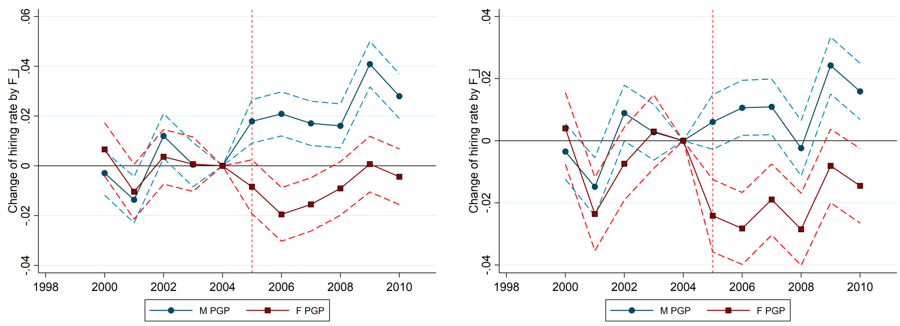

Note: This figure reports the regression coefficients capturing the effects of eliminating stated gender preferences on the hiring of females. Coefficients of the interaction term between year and indicators for job ads classified as advertising for women ("Female PGP"), or for men ("Male PGP") are reported. Dotted lines show the 95% confidence intervals. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.A: classification uses regression, B: classification uses random forests.

Source: AMS-ASSD data, own calculations.

Next we turn to difference of difference estimates derived from equation (1). As noted we use the 2000-2004 data as the "pre" period and 2006-2010 as the "post" period, removing the transition year 2005 from the analysis. Figure 2.7 graphically summarizes the coefficients from equation (1). We report the main effects of the PGP variables ($\beta_1$and β_2) in the upper part of the figure. As a point of reference we also show a 0 coefficient representing the comparison group of vacancies with neither PGP.

We report the interactions between the PGP dummies and the post-period dummy ($\lambda_1$and λ_2) in the lower part of the figure, again showing a 0 coefficient for the comparison group of vacancies with neither PGP. The left-hand figure (panel A) shows results based on our LOM prediction model, while the right-hand figure (panel B) shows results based on our RF model. (The main coefficients from these models (and all subsequent models) are presented in Appendix 2.E: tables 2.E1, 2.E2, and 2.E3.)

Figure 2.7: Female Hiring (Difference in Differences Results)

A. Leave-Out-Mean B. Random Forest

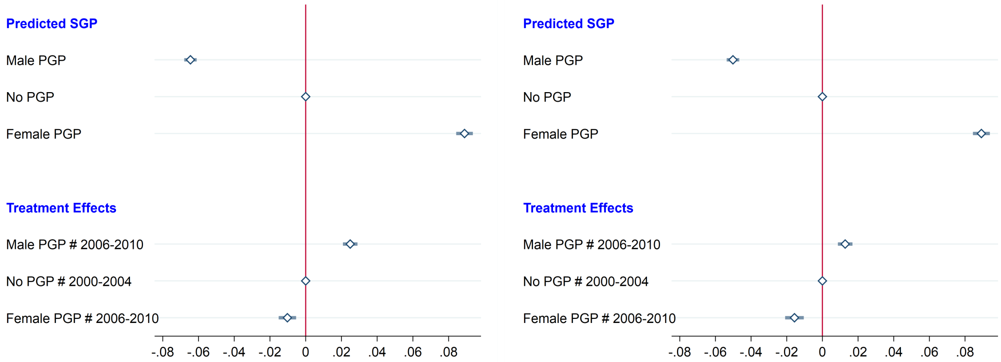

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females. The regressions follow the model in equation 2.3.1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects. A: classification uses regression, B: classification uses random forests.

Source: AMS-ASSD data, own calculations.

Using the LOM classifier, in the pre period vacancies with a male PGP have a 6.5 percentage point (ppt) lower probability of hiring a female candidate relative to vacancies with neither PGP, while vacancies with a female PGP have an 8.9 ppt higher probability of hiring a female. The estimates from the RF prediction model are similar, but show a slightly smaller impact of a male PGP (-5.5 ppts) and a slightly larger impact of a female PGP (10 ppt).

Note that if we assume an 80% attenuation in the effects of PGP status relative to true SGP status, then these estimates imply that stating a male SGP leads to 35 ppt lower probability of hiring a female relative to no SGP, while stating a female SGP leads to a 45 ppt increase. The implied gap in the probability of female hiring between female and male SGP's is therefore around 80 ppt, consistent with the patterns observed in Figures 3 and 4.

Looking next at the treatment effects, the estimates in Panel A (based on the leave-out-mean classifier) show that eliminating gender preferences increased the hiring of women by 2.5 percentage points for vacancies with a male PGP (standard error = 0.002), which is about 40% of the pre-campaign gap in the female hiring rate relative to vacancies with neither PGP. The estimated effect on the probability of hiring a female for vacancies with a predicted female PGP is -1.0 ppt (s.e. = 0.002), which is about 12% of the pre-campaign gap for these vacancies relative to the comparison group. RF estimates in Panel B suggest a slightly smaller gain (+1.6 ppts) in the hiring of females for vacancies with a male PGP, and a slightly larger reduction in the hiring of females (-2.7 ppts) for vacancies with a female PGP, consistent with the patterns in Figure 2.7.

Accounting for the attenuation arising from the slippage in our prediction models, these implied treatment effects are large in magnitude. The LOM-based estimates suggest that the elimination of gender preferences in job ads led to a roughly 15 ppt increase in the probability of hiring women to fill jobs that would have been advertised with a male gender preference in the absence of the new law, and 5 ppt increase in the hiring of men to fill jobs that would have been advertised with a male gender preference. The combination of the two effects therefore closed about one quarter (=(15+5)/80) of the previous gap in the rate of hiring women between jobs that specified a gender preference for men versus women.

As noted in the discussion of Table 1, vacancies filled by job searchers who had obtained services from the AMS differ somewhat from the overall population of vacancies. To assess the possible impact of these differences on our main models, we developed a set of weights reflecting the inverse probability that a given vacancy was filled by an AMS client, and re-estimated the models in Figure 2.7. Reassuringly, the results, reported in Appendix Figure 2.D1, are qualitatively very similar to those in Figure 2.7.

HETEROGENEITY BY FIRM SIZE Gender based hiring might differ between small and large firms. To probe the sensitivity of our results to firm size, we exclude firms that had less than five employees during the year before the vacancy was posted. For simplicity, in this analysis and all subsequent analysis we limit attention to models based on the LOM classifier. Figure 2.8 summarizes estimation results for our difference in differences model applied to larger firms. The pattern of results is very similar to the baseline pattern (Figure 2.6), suggesting that our results are not much affected by including very small firms.

Figure 2.8: Results for Large Firms

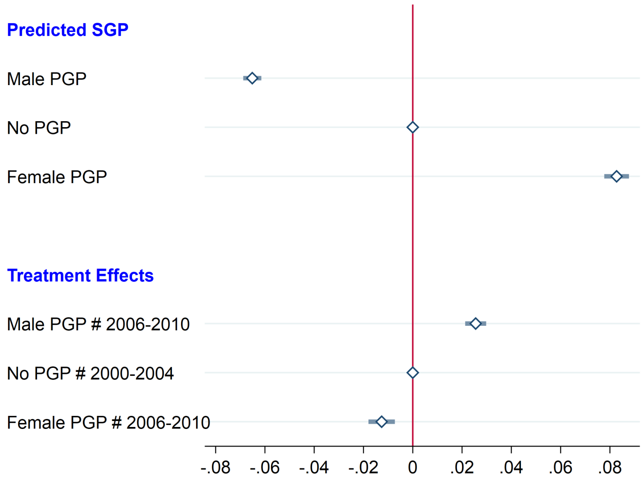

Note: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females for firms with more than 5 employees. The regression follows the model in equation 2.3.1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.
Source: AMS-ASSD data, own calculations.

VIENNA VS. THE REST OF AUSTRIA: Vienna is by far the largest labor market in Austria, with a relatively high share of employment in the services sector and a lower share in manufacturing and other more traditional sectors. We therefore estimated our baseline model separately for Vienna and the rest of Austria, with results in Figure 2.9.

Looking at the pre-campaign differences in the upper part of the figure, we see that vacancies with a male PGP had only about a 2.5 ppt lower probability of hiring a female in Vienna relative to vacancies in the comparison group, versus a roughly 7 ppt gap in the rest of Austria. In contrast, for vacancies with a female PGP the pre-campaign gap relative to the comparison group was about 10 ppt in Vienna versus 9 ppt in the rest of Austria.

After the campaign the probability of hiring a female for male PGP vacancies increased by about 1 ppt in Vienna – closing about 1/3 of the previous gap – while in the rest of the country it rose by about 2.5 ppt – again closing about 1/3 of the gap. For female PGP's the probability of hiring a female fell by about 3 ppt in Vienna (nearly one quarter of the pre-campaign gap) while in the rest of the country the decline was smaller (about a 1 ppt reduction, closing only 1/10 of the pre-campaign gap relative to the comparison group).

Figure 2.9: Vienna vs the rest

(a) Vienna (b) the rest of Austria

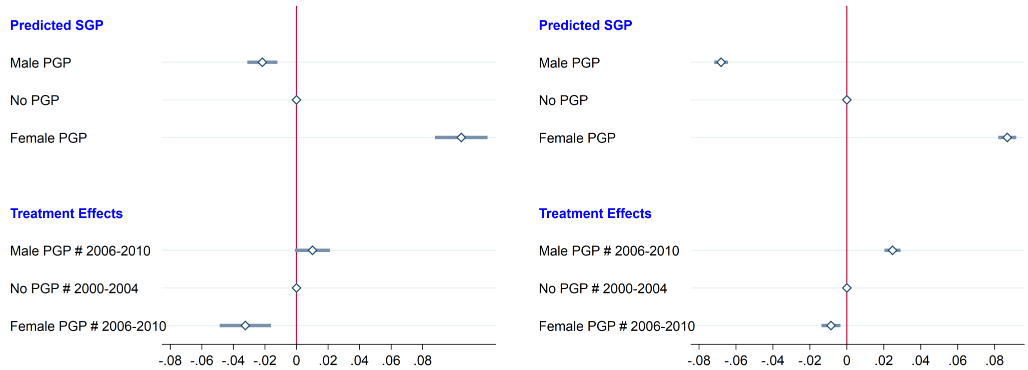

Note: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females for firms in Vienna (panel A) and in the rest of Austria (panel B). The regression follows the model in equation 2.3.1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

These results suggest that in the pre-campaign period the impacts of SGP's varied somewhat across labor markets, but that in all cases the elimination of posted gender preferences led to some gender diversification in hiring. The effects were particularly large for jobs with male PGP's outside of Vienna, and for jobs with female PGP's in Vienna. We speculate that more heavily male blue collar jobs were particularly segregated in the smaller towns and rural areas of Austria, and opened up more after the campaign, while more heavily female pink collar jobs in services and government were more segregated in Vienna, and also opened up more after the campaign.[21]

[21]Another concern with our baseline analysis is the presence of a fair amount of seasonal work in Austria. In a sensitivity analysis, we remove vacancies that offer seasonal jobs (12.1% in 2000-2004, 11.8% in 2006-2010), and find the same pattern of results as in our baseline analysis. Results are available upon request.

STEREOTYPICAL VS NON-STEREOTYPICAL VACANCIES

Our baseline difference of differences specification makes no distinction between SGP's that specify a preference for a gender that matches the majority of the existing workforce of the firm (stereotypical preferences) or for the opposite gender (non-stereotypical preferences). Our simple conceptual model of SGP's, however, suggests that vacancies with a predicted non-stereotypical gender preference could be most affected by the elimination of gender preferences. Moreover, the contrast between stereotypical and non-stereotypical preferences is interesting because the former help maintain gender segregation, while the latter promote gender diversity. In this section we expand our framework and explore heterogeneity along this dimension.

To proceed we label workplaces with more than 50% females as mainly female, indicated by a dummy variable $C_j^f = 1[F_j >= 0.5]$, and those with less than 50% females as mainly male, indicated by a dummy variable $C_j^m = 1[F_j < 0.5]$.[22] A vacancy with a given predicted SGP is stereotypical if the predicted gender preference is concordant with the firm's existing gender composition ($D_j^f C_j^f = 1$, stereotypical female; or $D_j^m C_j^m = 1$, stereotypical male) and non-stereotypical if the predicted gender preference is discordant with the existing workforce composition ($D_j^f C_j^m = 1$, non-stereotypical female; or $D_j^m C_j^f = 1$, non-stereotypical male).

Table 2.5 presents descriptive information on stereotypical and non-stereotypical vacancies in the pre-campaign period (where we can see actual gender preferences as well as our predicted preferences). Columns 1 and 2 show characteristics of vacancies with a male ***predicted*** gender preference (PGP), classified further by whether the employer had a mainly female workplace (column 1) or a mainly male workplace (column 2). Note that only 7% of male PGP vacancies (in column 1) are non-stereotypical. Similarly, columns 4 and 5 show characteristics of vacancies with a female PGP, classified by whether the firm has a majority female workplace (column 4 - the stereotypical case) or a mainly male workplace (column 5 – the non-stereotypical case). Again, only a small share (12.8%) of female PGP vacancies are non-stereotypical.

Panel A of the table shows the shares of each PGP that had actual stated preferences for males, females, or neither. Interestingly, we see that the fraction of concor-

[22]This classification is the same as we used on Table 2.3 to examine vacancy filling times in the pre-campaign years.

Table 2.5: Descriptive Statistics for Stereotypical and Non-Stereotypical Predicted Gender Preferences

	Predicted Gender Preference				
	M PGP in F WP [Non-stereo Male PGP]	**M PGP in M WP** [Stereo Male PGP]	**Neutral**	**F PGP in F WP** [Stereo Female PGP]	**F PGP in M WP** [Non-stereo Female PGP]
	(1)	(2)	(3)	(4)	(5)
Panel A: SGP					
Male SGP	0.5881	0.6073	0.1820	0.0168	0.0617
Neutral	0.3719	0.3802	0.6314	0.4387	0.4565
Female SGP	0.0400	0.0126	0.1867	0.5445	0.4818
Panel B: Context					
Share of women in the firm	0.6349	0.1330	0.4696	0.8810	0.3068
Share of women in occupation	0.3728	0.1952	0.5203	0.7075	0.6703
Observations	3,676	51,841	114,273	41,745	6,033
Shares	0.0169	0.2383	0.5252	0.1919	0.0277

This table shows summary statistics for stereotypical (Male PGP in Male Workplace in column (2) and Female PGP in Female Workplace column (3)) and non-stereotypical (Male PGP in Female Workplace in column (1) and Female PGP in Male Workplace column (4)) PGP's. Period 2000-2004.

dant actual and predicted gender preferences is about the same whether the PGP is stereotypical or not. For example, 58.81% of stereotypical male PGP's have an actual male SGP (column 1), versus 60.73% of non-stereotypical male PGP's. The gap is only slightly wider between stereotypical and non-stereotypical female PGP's (54.45% versus 48.18%).

Panel B shows the shares of female workers at the posting firms and in the relevant occupations for each PGP group. By construction stereotypical and non-stereotypical vacancies come from firms with much different shares of female workers. The differences in gender shares in the recruited occupations, however, are much smaller. For example, a typical occupation for a non-stereotypical female PGP vacancy (column 5) was 67.03% female, whereas a typical occupation for a stereotypical female PGP vacancy (column 4) was 70.75% female. The gap is somewhat larger for male PGP vacancies (~18 ppt) but is still only about 1/3 as large as the gap in the share of women at the firm's workplace (50 ppts).

We use an extended version of equation (1) to measure the effects of the 2004 campaign on 4 distinct types of vacancies based on PGP and workplace gender composition:

$$\begin{aligned} y_j &= \beta_0 + \beta_1 D_j^f C_j^f + \beta_2 D_j^f C_j^m + \beta_3 D_j^m C_j^m + \beta_4 D_j^m C_j^f \\ &+ \lambda_1 D_j^f C_j^f Post_j + \lambda_2 D_j^f C_j^m Post_j + \lambda_3 D_j^m C_j^m Post_j + \lambda_4 D_j^m C_j^f Post_j + X_j\delta + \varepsilon_j \end{aligned} \quad (2.5.1)$$

In this specification the β_k coefficients measure differences across vacancies in the outcome y before the campaign is implemented, while the λ_k coefficients measure changes in the the outcome after the introduction of the campaign (λ_1 and λ_3 for stereotypical vacancies, and λ_2 and λ_4 for non-stereotypical vacancies). As in the simpler model (1), the λ_k coefficients are attenuated by misclassification errors relative to the coefficients in a model that relates the outcome to the stated gender or desired gender preference. In the Appendix we present an analysis similar to that in equations (3)-(6) discussing the attenuation effects. As shown in columns 2 and 4 of Table 2.4, when we relate actual SGP's to predicted gender preferences, allowing different effects for stereotypical and non-stereotypical PGP's, we find that non-stereotypical PGP's of a given gender are stronger predictors of actual SGP's of that gender, but are also more negatively related to the SGP of the opposite gender. Allowing for the multiple channels we conclude that the expected attenuation of the effects of stereotypical and non-stereotypical vacancies in a model like equation (7) are likely to be similar, and in both cases close to a factor of 80%.

Estimation results for equation (7) are summarized in Figure 2.10.[23] The estimates show that in the pre-campaign period non-stereotypical PGP's had relatively large effects on the probability of hiring a female candidate (as would be expected from the "Z"-shaped pattern in Figures 3-4), whereas stereotypical PGP's had more modest effects, comparable to the overall effects from specification (1) reported in Figure 2.10. Specifically, females were 4.7 ppts less likely to be hired for vacancies with a stereotypical male PGP (relative to the omitted group of vacancies with no preference), but 18.7 ppts less likely to be hired for vacancies with a stereotypical male PGP. Conversely, females were were 5.7 ppts more likely to be hired for vacancies with a stereotypical female PGP and 19.3 ppts more likely to be hired for vacancies with a stereotypical female PGP.

Likewise, the difference in differences estimates show bigger effects of the elimination of SGP's on the hiring outcomes of vacancies with non-stereotypical PGP's: a 6.8 ppt increase (s.e.=0.008) in the probability of hiring a female for a non-stereotypical male PGP versus a 2.3 ppt increase (s.e.=0.002) for a stereotypical male PGP; and a 4.9

[23] As in our previous models we include controls for time, industry (86 dummies), occupation (100 dummies), and the proportion of females at the workplace in the year prior to the posted vacancy (51 dummies).

Figure 2.10: Effects on Female Hiring (Stereotypical vs Non-Stereotypical Vacancies)

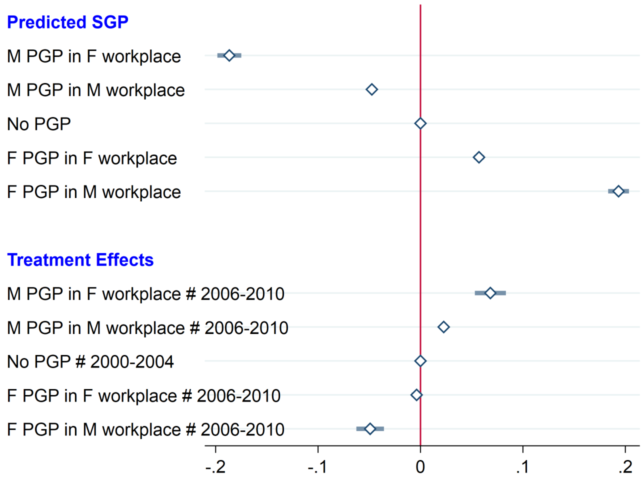

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females for Stereotypical (Predicted M SGP in M Workplace and F SGP in F Workplace) and Non-Stereotypical (Predicted M SGP in F Workplace and F SGP in M Workplace) vacancies. The regression follows the model in equation 2.5.1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

ppt reduction (s.e.=0.007) in the probability of hiring a female for vacancies with a non-stereotypical female PGP's, versus an insignificant 0.004 ppt reduction for vacancies with a stereotypical female PGP.

These estimates imply that most of the overall increase in the hiring of females for vacancies with a male PGP was attributable to gains at majority-male workplaces (which tended to use stereotypical male gender preferences in the pre-campaign period). Interestingly, there was no corresponding increase in the hiring of males for vacancies with a female PGP at majority-female workplaces. Instead, the gains in hiring for men were driven entirely by a rise in the fraction of vacancies at mainly male workplaces with a predicted female preference that were filled by men once SGP's were eliminated.

Effects on Vacancy Durations and Other Job Outcomes

The elimination of gender preferences in recruiting could affect both the cost of forming new job matches and the quality of these matches. To assess these effects we examine three outcomes: the number of days required to fill a vacancy ("fill time"); the starting wage at the newly created job, and the duration of the newly created job. We interpret fill time as a measure of the cost of forming a match and the duration of the newly created job as a measure of the quality of the match. The interpretation of wage outcomes is less clear, though as we will see impacts of predicted gender preferences and the elimination of such preferences on wages are small. We conduct our analysis in the framework of equations (1) and (7), simply replacing the dependent variable with the match outcome measures. All estimates condition on the three sets of fixed effects included in our previous models (industry, occupation, and lagged female share) as well as year fixed effects.

Looking first at vacancy fill times, the results in Panel A of Figure 2.11, show that in the pre-campaign period vacancies with a female PGP had shorter fill times. The elimination of SGP's had little or no effect on filling times for vacancies with male PGP, but significantly increased the time to fill those with a female PGP. The estimates in Panel B, based on specification (7) provide some additional insights. In particular, we see that average fill times for both stereotypical and non-stereotypical male PGP's were about the same in the post-campaign period as before. For vacancies with a stereotypical female PGP we also see no significant change in fill times. All of the rise in average fill times for predicted female PGP's is attributable to an increase in vacancy durations for non-stereotypical female PGP's.

The lack of effects on vacancy filling times for stereotypical male PGP's is interesting. Contrary to what might have been expected given the assumptions of Kuhn and Shen (2013), we do not see a slow down in the vacancy filling rate caused by "congestion" in the candidate assessment process for such vacancies. Taken together with the fact that more female candidates were hired to fill these vacancies in the post-campaign period, the evidence suggests that some employers had been using invalid or out-of-date priors to limit their applicant pools. Once SGP's were eliminated, they learned that there were in fact acceptable female candidates and actually hired some of them, taking about the same time to reach a decision with a presumably larger applicant pool.

The most notable post-campaign effect in Panel B of Figure 2.12 is the approximately 2 day rise in filling times for vacancies with non-stereotypical female PGP's (a roughly 7% increase). We saw in Figure 2.10 they these openings were more likely to be filled by men, suggesting that some employers with a majority male workforce that

Figure 2.11: Effects on Vacancy Filling

A. Baseline Model B. Stereotypical vs Non-Stereotypical Vacancies

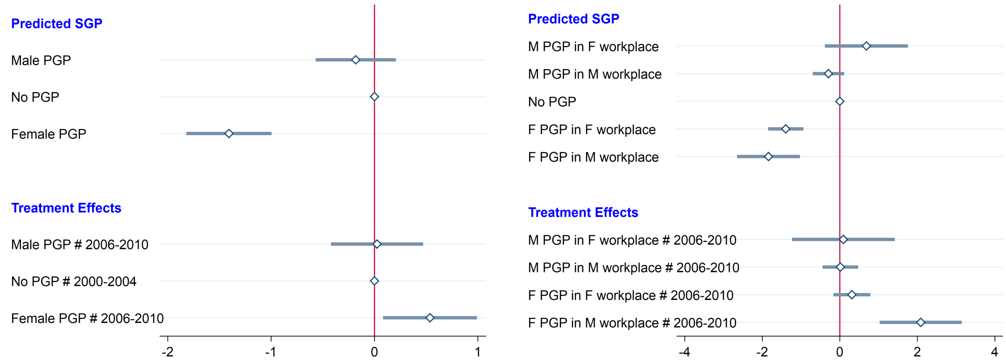

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females for Stereotypical (Predicted M SGP in M Workplace and F SGP in F Workplace) and Non-Stereotypical (Predicted M SGP in F Workplace and F SGP in M Workplace) vacancies. The regression follows the model in equation 2.5.1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

*Source:*AMS-ASSD data, own calculations.

wanted to recruit a female worker experienced a reduction in applications from females in the post-2005 period and ultimately had to compromise on a male candidate, with some overall slow down in the decision process.

Next we turn to job durations and wages. In our main models for these outcomes we add controls for the gender, age and previous wage of the new recruit. While the characteristics of the newly hired worker are endogenous, we believe it is most straightforward to interpret the impacts of the elimination of SGP's on wages and job durations with these controls in place, since female workers have lower wages and longer job durations across all jobs, and a change in the gender of the new recruit would be expected to mechanically affect these outcomes. We present models without the extra controls in Appendix Figure 2.D3, and discuss the differences relative to models with these controls below.

Figure 2.12 summarizes the results for job durations. An issue in constructing the dependent variable for this analysis is right censoring of jobs that were still in progress at the time we extracted our sample. We address this by censoring all jobs at four years (1440 days). Since only about 10 percent of jobs last more than four years, we believe this censoring has little impact on our findings.

Figure 2.12: Effects on Job Duration

C. Baseline Model D. Stereotypical vs Non-Stereotypical Vacancies

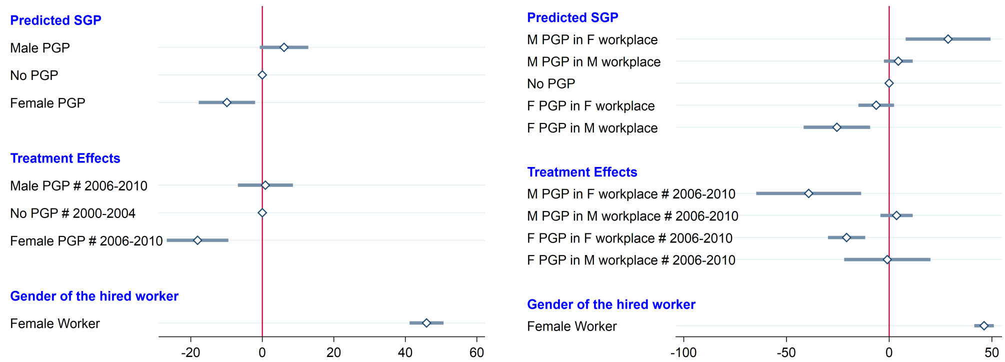

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on completed job durations. In panel A the the model follows the baseline specification in equation 2.3.1. In panel B the model follows the specification in equation 2.5.1, distinguishing between Stereotypical (M PGP in M Workplace and F PGP in F Workplace) and Non-Stereotypical (M PGP in F Workplace and F PGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include previous wage, gender, age, as well as industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations. calculations

The estimates in panel A of Figure 2.12 show that in the pre-2005 period durations were insignificantly longer for jobs created by vacancies with a male PGP (relative to those with no predicted preference) but significantly shorter (by about 19 days) for jobs that filled vacancies with a female PGP. We also see (at the bottom of the figure) that jobs filled by women tend to last considerably longer (by about 45 days) conditional on industry, occupation and the gender composition of the employer's workforce – consistent with the findings in Table 2. In the post-campaign period there was no change in the average duration of jobs created by vacancies with a male PGP but a further $\sim$ 14 day reduction in the duration of jobs created by vacancies with a female PGP, equivalent to a roughly 7 percent effect relative to the median.

The results in Panel B, based on equation (7), provide some further insights into these treatment effects. In particular, we we can see that the relatively constant average duration of jobs arising from vacancies with a male PGP reflects a combination of a small (insignificant) increase in the duration of jobs associated with stereotypical male PGP's and a large reduction in the duration of jobs associated with non-stereotypical male PGP's. The decline in the average duration of jobs from vacancies with a female PGP, on the other hand, is entirely attributable to a fall in the duration of jobs associated with stereotypical female PGP's (i.e., jobs at mainly female workplaces).

We suspect that the shortened duration of jobs associated with non-stereotypical male PGP's is reflective of a reduction in match quality for these jobs (which are more likely to be filled by female candidates in the post-campaign period), coupled with a general reduction in job durations at mainly female workplaces that also affects the jobs associated with stereotypical female PGP vacancies. To examine this further, we conducted an event study analysis of job durations associated with male and female PGP's (see Appendix Figure 2.D5). A visual examination suggests that durations of jobs associated with female PGP's (the vast majority of which are at mainly female workplaces) were trending downward by 2-3 days per year throughout our sample period. This trend can account for most of the apparent treatment effect for stereotypical female PGP vacancies.

The elimination of stated gender preferences might also affect the wages offered to new recruits (conditional on gender, age, and past wage). As shown in panel A of Figure 2.13 we find that, prior to the campaign, daily earnings were about 3 percentage points lower for jobs associated with vacancies with a female PGP, and around 1 percentage point higher than in the comparison group for jobs arising out of vacancies with a male PGP. Neither set of jobs experience significant changes in pay after the campaign. Looking in panel B of the figure, we see that in the pre-2005 period there was a significant wage disadvantage for non-stereotypical male PGP jobs (mostly jobs held by men at mainly female workplaces) as well as for jobs associated with both stereotypical and non-stereotypical female PGP's. In the post-campaign period there was a modest, marginally significant gain in wages for non-stereotypical male PGP jobs but no large change for any other type of new job. Indeed, apart from the gain in wages for non-stereotypical male PGP jobs, we can rule out wage changes that are bigger than ± 2 percentage points at conventional levels of statistical significance.[24]

It is worth emphasizing that the models in Figure 13 control for the gender age and previous wage of the new recruits. When we exclude these controls, we find that mean

[24]There are some differences in the part-time versus full time status of male versus female PGP's: vacancies with a female PGP tend to have a larger share of part time work, while those with a male PGP have a smaller share of part time work.

Figure 2.13: Effects on Wages

A. Baseline Model

B. Stereotypical vs Non-Stereotypical Vacancies

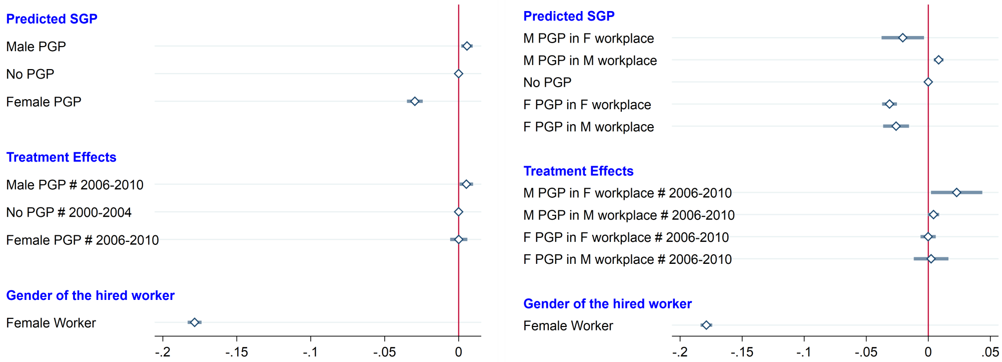

Notes: This figure reports estimation results capturing the effect of eliminating stated gender preferences on wages for newly filled jobs. In panel A the the model follows the baseline specification in equation 2.3.1. In panel B the model follows the specification in equation 2.5.1, distinguishing between Stereotypical (M PGP in M Workplace and F PGP in F Workplace) and Non-Stereotypical (M PGP in F Workplace and F PGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include previous wage, gender, age, as well as industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations. calculations

wages for jobs associated with non-stereotypical male PGP's fell in the post-campaign period while wages for jobs associated with non-stereotypical female PGP's rose. We interpret both of these impacts as a result of compositional effects: in the former case because of a rise in the share of women hired to fill vacancies with non-stereotypical male PGP's; in the latter because of a rise in the share of men hired to fill vacancies with non-stereotypical female PGP's.

2.4.3 Effects on Workplace Diversity

As a final exercise, we return to the issue of the gender of the newly hired work and ask explicitly how the new recruit affected the gender diversity of the employer's workforce. To quantify this effect we define a simple measure $d_j \equiv |H_j - F_j|$, where H_j is a dummy equal to 1 if the new hire for vacancy j is female, and F_j represents the share of female workers at the workplace posting the vacancy (measured in the previous year). This indicator, which ranges from 0 (a completely segregated workplace hires a new worker of the same gender as all existing workers) to 1 (a completely segregated workplace hires a new worker of the opposite gender), summarizes the deviation of the gender of the new hire from the composition of the workforce. The index will be higher if new hires are more likely to work with co-workers of the opposite gender.[25] The mean value of the index during the pre-2005 period was about 0.22.

Figure 2.14 summarizes the estimates from our simpler specification (1) (in Panel A) and from our richer specification (7) that distinguishes between stereotypical and non-stereotypical PGP's (in Panel B). Looking first at Panel A, we see that in the pre-campaign period, use of SGP's was associated with reductions in diversity, relative to the comparison group of vacancies with neither PGP. This reflects the fact that, e.g. vacancies with male PGP recruit men into male workplaces, and similarly for women. The effect of eliminating SGP's was to increase the probability of a diversity-enhancing hire for both male and female PGP vacancies, with an effect of +0.007 (s.e.=0.001) for male PGP's and +0.006 (s.e.=0.002) for female PGP's. Recalling that these effects are attenuated by a factor of roughly 80% relative to a specification using stated/desired SGP's, the estimates suggest relatively large increases in the diversity of new hires (+.035 for male male PGP's and +0.030 for female SGP's) relative to a pre-campaign mean of 0.22. Stated differently, we estimate that the campaign to eliminate stated gender preferences reduced the tendency to hire non-diversifying candidates at workplaces that would use these preferences by a factor of 1/6. Of course many other factors that contribute to non-diverse hiring (such as location of the workplace, schedule of work, and attitudes of managers and coworkers) were unaffected by the campaign.

[25]Specifically, rearranging the indicator shows $E(d_j) = E(M_j|H_j = 1)P(H_j = 1) + E(F_j|H_j = 0)P(H_j = 0)$, where $M_j \equiv 1 - F_j$ is the share of males at the workplace. The first term in this expression is the share of men at workplaces hiring women times the share of new hires that are women; the second term is the share of women at workplaces hiring men times the share of new hires that are men. Holding constant the gender composition of all new hires, this indicator will rise if the coworkers of new hires are more likely to be of the opposite gender. A second way to measure diversity is, $e_j \equiv H_j(1 - C_j^f) + (1 - H_j)C_j^f$ where C_j^f identifies female firms. This second index identifies female hires in male firms, $H_j(1-C_j^f)$, and male hires in female firms, $(1 - H_j)C_j^f$. This index bears a relationship with the Duncan index of gender segregation index (Duncan and Duncan, 1955), $D_j = \frac{1}{2}\sum_j |\frac{m_j}{M} - \frac{f_j}{F}|$, where m_j is the number of men, f_j is the number of women employed at workplace j, and M and F are the total number of men and women. Suppose the total numbers of men and women are identical, e.g. $M = F$, then the Duncan index then is $D_j = \frac{1}{2F}\sum_j[(m_j - f_j)(1 - C_j^f) + (f_j - m_j)C_j^f]$, capturing by how much men outnumber women in male firms, and women outnumber men in female firms, on average. Suppose we compare segregation over two time periods in a situation where firms hire only one person, and no person leaves the firm. The change in segregation across these two time periods is (up to normalization) $\Delta D_j \propto \sum_j[(\frac{1}{2} - H_j)(1 - C_j^f) + (H_j - \frac{1}{2})C_j^f] = \sum_j \frac{1}{2}[(1 - C_j^f) - C_j^f - e_j]$. In the appendix (Figure 2.D6), we report results for $\frac{1}{N}\sum_j e_j$, and these results can be interpreted as the reduction in segregation, as measured by the Duncan index, induced by the campaign. Results are consistent across both types of indices.

Figure 2.14: Effects on Workforce diversity

A. Baseline Model B. Stereotypical vs Non-Stereotypical Vacancies

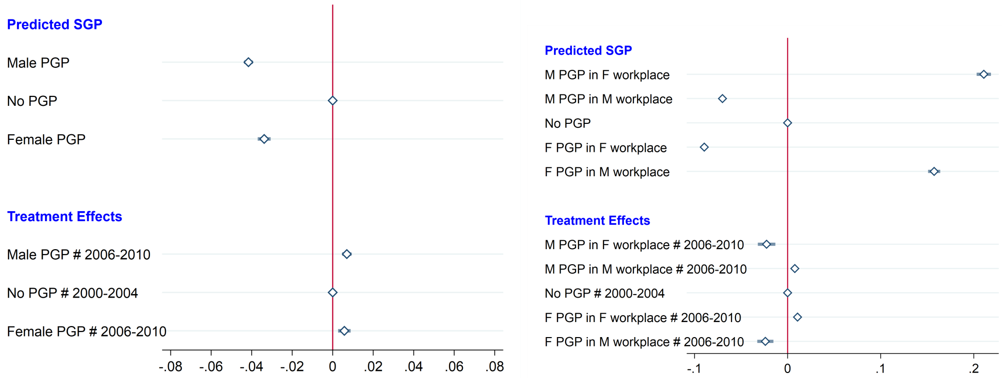

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on workplace composition. In panel A the model follows the baseline specification in equation 2.3.1. In panel B the model regression follows the specification in equation 2.5.1, distinguishing between Stereotypical (Predicted M SGP in M Workplace and F SGP in F Workplace) and Non-Stereotypical (Predicted M SGP in F Workplace and F SGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

Looking at the estimates from our richer specification (Panel B), we see a sharp contrast between stereotypical and non-stereotypical PGP's in both the pre-period and the post-campaign period. When stated SGP's were allowed, our estimates suggest that non-stereotypical SGP's substantially increased the probability of a diversity-enhancing hire, while stereotypical SGP's had the opposite effect. The effect of eliminating stated gender preferences was to increase the probability of a diversity-enhancing hire for vacancies that are predicted to have had a stereotypical gender preference, but to reduce the probability of a diversity-enhancing hire for those predicted to have had a non-stereotypical gender preference.[26] These patterns are consistent with the effects of the campaign on the gender of the candidates hired to fill stereotypical and non-stereotypical PGP's noted in Figure 10. But since stereotypical vacancies account for 95% of all vacancies, the increase in diversity on stereotypical hires dominates the loss in diversity on non-stereotypical vacancies. Interestingly, we also see that non-

[26]The result on stereotypical female vacancies is noteworthy. Firms that posted stereotypical female vacancies did not hire significantly more men (Figure 2.10B), but the male and female hires entered more workplaces with many co-workers of the opposite gender. Abolishing stereotypical job ads can increase diversity even without affecting hires directly.

stereotypical vacancies continue to reduce segregation to an important extent, even after the campaign.

2.5 Conclusion

In early 2005 the *Ombud for Equal Treatment* (OET) agency in Austria began a publicity campaign to inform employers and newspapers that gender preferences in job advertising were illegal. Over the next year the use of stated gender preferences on the job board of the Austrian Employment Service fell from around 40% of all vacancies to less than 5%. We use data on filled vacancies from this board to study how the elimination of stated gender preferences affected hiring outcomes and job quality measures. To focus on affected firms, we use pre-2005 data on the occupation of the job opening, the firm's industry, and its lagged female employment share to predict the gender preferences (female, male, or neither) in a given vacancy. We further classify predicted preferences into stereotypical or non-stereotypical, based on the concordance between the predicted preference and the majority gender at the firm's workplace. Then we estimate simple difference in differences models, comparing pre- and post-campaign hiring outcomes by predicted gender preference group. These models can be interpreted as reduced form estimates from a system that expresses hiring outcomes in terms of latent gender preferences that are only observed in the pre-2005 period.

Prior to the OET's campaign we find that most stated gender preferences (90% or more) were aligned with the gender composition of the recruiting firm's existing workforce. Thus, on average, stated gender preferences tended to reinforce existing patterns of gender segregation. Nevertheless, a small fraction of employers posted non-stereotypical preferences, requesting candidates that were the opposite gender of the majority of their workforce. Consistent with a series of papers looking at call-back rates to job applicants (Kuhn and Shen, 2013; Helleseter et al., 2018; Kuhn et al., 2020) we find that stated gender preferences were strongly predictive of actual hiring outcomes. We also find that vacancies with both stereotypical and non-stereotypical gender preferences had faster filling times, suggesting that gender preferences served as reliable and salient signals to job searchers, and that there were many workers available to fill even non-stereotypical job openings.

Our difference in differences models show that the elimination of stated gender preferences led to a significant 2.5 percentage point increase in the fraction of women hired to fill vacancies with a predicted male preference, and a smaller (but still significant) 1 point increase in the fraction of men hired to fill vacancies with a predicted female preference. Scaling these effects to reflect the attenuation between predicted and actual gender preferences we infer that the OET campaign had relatively large effects on job opportunities, particularly for women.

Looking further into stereotypical versus non-stereotypical preferences, we find that the gains for women were driven by a rise in hiring to fill job openings at majority-male workplaces with predicted male preferences (i.e., stereotypical male preferences). There was also an increase in the hiring of women at majority-female workplaces that were predicted to use male preferences, offset by a reduction in hiring at majority male workplaces that were predicted to use female preferences. The rise in female hiring for openings with stereotypical male preferences implies a reduction in gender segregation across workplaces, while the shifts in hiring for vacancies with non-stereotypical preferences work in the opposite direction. Given the relatively small shares of non-stereotypical preferences, however, the net effect was still toward desegregation. Examining segregation effects directly using a variant of the Duncan index, we calculate that eliminating gender preferences reduced gender segregation by about 16% at workplaces that would have used gender preference statements.

Looking at filling times, wages and durations of the newly created jobs, we find little evidence that the elimination of stated preferences led to a reduction in the speed of matching or the quality of job matches for the large group of vacancies with stereotypical gender preferences. For vacancies that were predicted to have non-stereotypical preferences we see some increase in recruiting times of openings with a predicted preference for men in female firms, and declines in the duration of jobs of openings with a predicted preference for men at a mainly-female workplace. Such effects are not necessarily surprising, given the relative rarity of such vacancies and the difficulty in signaling the firm's intended hiring target given other information about the job and the workplace.

This pattern of evidence suggest that both job seekers and firms adapted to the removal of gender preference statements after the OET campaign. On the worker side, our findings on hiring outcomes, and related findings from Kuhn and Shen (2021) on the responses of job applicants to the removal of SGP's, suggest that applicant pools became more gender-diverse after the campaign. In our setting this change may have been driven in part by changes in the referrals made by Employment Service caseworkers. On the firm side, our most surprising finding (again, consistent with findings reported by Kuhn and Shen, 2021 based on employer call-backs) is that some firms that would have specified a gender preference prior to the campaign responded to the newly diverse applicant pool by actually hiring workers of the opposite gender, rather than simply ignoring their applications. In most cases the net effect of these responses was to increase the gender diversity of hiring, though in cases where an employer would have used a non-stereotypical gender preference the effect was to reduce gender diversity of hiring.

Overall, our findings suggest that any negative effects of eliminating early job market signals of gender preference were mainly confined to the set of firms with non-stereotypical hiring goals. For the large set of firms with stereotypical male preferences, eliminating the ability to advertise these preferences led to an increase in hiring of women. Taken together with the finding of no degradation in match quality for the associated jobs, we conclude that many of the stereotypical preferences observed before the campaign were likely based on outdated priors, rather than on true productivity gaps or on rigidly held discriminatory beliefs that would be immune to policy.

APPENDIX

2.A DATABASE CONSTRUCTION

The AMS vacancy database contains the stock of all the open vacancies with a monthly frequency. AMS also records the inflow (posting date) and the outflow (closing date), as well as the outcome of each vacancy. Each vacancy can result with (i) the hiring of a worker through direct mediation of the AMS system, (ii) the hiring of a worker through different channels or (iii) no hiring (or no information about the hiring)[27].

The full database contains about 13.9 millions observations for the period 1997-2013, among which 5.2 millions are recorded as outflows. We consider only vacancy outflows, and first step toward the construction of our database is providing an opening date for each vacancy in the outflow subsample. AMS record a vacancy identifier, however, the identifier is not always reliable because an previously used identifier can be re utilized for a new opening in the future. We then match vacancy in time (inflows and outflows) by using the vacancy identifier (*vdgnr*) and other 10 time invariant vacancy specific variables. In this process the sample size reduces by 118,136 units (2%).

Our empirical analysis focuses only on vacancies filled through AMS, since they contain a person identifier (*penr*) that allows us to get information on the hired person, the gender in particular. This subsample contains 1.2 millions vacancies, corresponding to about one quarter of all the hires[28].

The second step of our database construction is matching the AMS vacancy outflows with the employment spells in the ASSD database (*qualifikation*). This is the most challenging part of the process since the firm identifiers in the AMS (*btrnr*) and in the ASSD (*benr*) database differ. We proceed as follows:

- We first "treat" the employment spells database. We normalize the spells by merging together all the spells between the same firm-worker pair with a break lower than 70 days and we drop all the spells related to social security institutions and to sick leave, unemployed or retired status, self employed and civil servants.

- We then match each individual hired in the vacancy database with the spell database through the worker id (*penr*). At this point for each vacancy we have a list of all

[27] This usually happens when the firm withdraw the job ad from the AMS because it is not interested in finding a new worker anymore or when the firm can not be contacted by the AMS any longer.

[28] The portion of lapsed vacancies is about 13% of the total number of outflows.

the employment spells of the hired worker and we need to identify the spell that the vacancy refers to. To do that we consider only the spells starting around the closing date of the vacancy (-40, +90 days) and we drop all spells ending before the closing date of the vacancy.

- At this point we have a unique vacancy-spell match for most of the vacancies. For the remaining ones, we compare the pair AMS firm identifier (*btrnr*) and ASSD firm identifier (*benr*) with the "*cross-walk*" document provided by the *Bundesministerium fuer Arbeit, Soziales und Konsumentenschutz*[29]. For vacancies with at least a match with the cross-walk document we get the first one in chronological order among the matched ones, for the others we get the first in chronological order among all the identified spells.

At the end of this procedure we are able to have a pair vacancy-employment spell for 88% of the sub-sample of vacancies for which we observe a worker identifier.

The last step is getting the firm information, the worker information and the earnings. We do this by matching our database with employers database and workers database through the firm id (*benr*) and the worker id (*penr*) contained in the employment spell data. The final *vacancy-employer-employee* sample contains 987,271 observations for the period 1997-2013.

Table 2.A1 shows how the number of observations decreases at each of the steps above, while tables 2.A2 and 2.A3 compare the distribution of industries and occupations, respectively, across the sub-samples.

[29]This document matches the two identifiers, even though we tested the validity of the crossing with poor results, we believe that it contains some information that we can exploit to refine our matching algorithm.

Table 2.A1: Observation by subsamples

	Observations
1. Vacancies from September 1997 to December 2013 (after restrictions) 13,906,275	
2. Outflows of Vacancies from September 1997 to December 2013	5,214,539
3. Outflows of Vacancies from Sept. 1997 to Dec. 2013 with Inflow date	5,096,403
(a) With non missing key variables (4,998,146)	
4. Hired by AMS from September 1997 to December 2013	1,169,203
5. Drop 414 duplicates	1,168,789
6. Merge with Employment Spells Database	1,042,794
7. Merge with Firms Database	1,041,094
8. Generate key variables, including occupations codes	1,016,843
9. Get worker personal information (gender, birthdate)	1,014,701
10. Get wages	987,271

Notes: This table shows the number of observations in the database at each stage of the sample construction for the period 1997-2013.

Source: AMS-ASSD data, own calculations.

Table 2.A2: Job Ads by Industry

Industry	**Subsample**			
	All Outflows (1)	**Hired** (2)	**Hired AMS** (3)	**Matched** (4)
Agriculture, hunting and forestry	0.24%	0.25%	0.17%	0.17%
Mining and quarrying	0.10%	0.10%	0.14%	0.15%
Manufacturing	10.47%	10.81%	14.16%	14.62%
Electricity, gas and water supply	0.09%	0.10%	0.09%	0.09%
Construction	7.23%	7.05%	8.73%	8.81%
Wholesale and retail trade	15.22%	15.82%	16.90%	17.37%
Hotels and restaurants	24.20%	25.39%	19.77%	19.26%
Transport, storage and communication	3.74%	3.87%	4.18%	4.26%
Financial intermediation	0.95%	0.83%	0.64%	0.65%
Real estate, renting and business activities	27.07%	25.44%	22.68%	21.92%
Public administration and defence; compulsory social security	1.99%	1.85%	2.38%	2.44%
Education	0.95%	0.86%	0.76%	0.77%
Health and social work	3.47%	3.47%	4.78%	4.88%
Other community, social and personal service activities	4.27%	4.14%	4.59%	4.60%

Notes: This table shows the share of job postings by broad categories of industries for four different subsamples. Period 1997-2013.

Source: AMS-ASSD data, own calculations.

Table 2.A3: Job Ads by Occupation

Occupation	**Subsample**			
	All Outflows (1)	**Hired** (2)	**Hired AMS** (3)	**Matched** (4)
Legislators, senior officials and managers	1.00%	0.99%	0.63%	0.65%
Professionals	1.96%	1.73%	0.87%	0.83%
Technicians and associate professionals	9.95%	9.31%	6.64%	6.82%
Clerks	6.79%	7.01%	8.20%	8.50%
Service workers and shop and market sales workers	33.89%	35.25%	30.52%	30.45%
Skilled agricultural and fishery workers	0.42%	0.42%	0.41%	0.41%
Craft and related trades workers	21.01%	19.88%	19.80%	20.06%
Plant and machine operators and assemblers	5.97%	6.06%	7.54%	7.67%
Elementary occupations	19.02%	19.34%	25.40%	24.61%

Notes: This table shows the share of job postings by broad categories of occupations for four different subsamples. Period 1997-2013.

Source: AMS-ASSD data, own calculations.

2.B ATTENUATION WITH (NON-)STEREOTYPICAL VACANCIES

We are using predicted SGP's rather than actual gender preferences, so equation (2.3.2) can be interpreted as the reduced-form from a two-stage-least-squares procedure in which we estimate the first stage models using only data from the pre-campaign period, and allow the effects of the endogenous variables to vary between the pre-campaign and post-campaign periods. To formalize this, let S_j^f and S_j^m represent dummies for actual SGP's in the pre-campaign period or *desired* SGP's in the post-campaign period (i.e., the preferences that employers would have stated if there was no effort to eliminate SGP's). Let C_j^f represent workplaces with a majority of women (F workplace), and C_j^f represents workplaces with a majority of men (M workplace). Assume that the true model generating outcome y is:

$$\begin{aligned} y_j = \alpha_0 + \alpha_1 S_j^f C_j^f + \alpha_2 S_j^f C_j^m + \alpha_3 S_j^m C_j^m + \alpha_4 S_j^m C_j^f \\ + \theta_1 S_j^f C_j^f Post_j + \theta_2 S_j^f C_j^m Post_j + \theta_3 S_j^m C_j^m Post_j + \theta_4 S_j^m C_j^f Post_j + X_j \gamma + \varepsilon_j \end{aligned} \tag{2.B0.1}$$

This has the same form as equation (2.3.2) but relates the outcome to true (or desired) SGP's. Assume that the actual/desired SGP's are related to predicted preferences by a pair of simple models with constant coefficients between the pre- and post-campaign periods:

$$S_j^f = \pi_0 + \pi_1 D_j^f C_j^f + \pi_2 D_j^f C_j^m + \pi_3 D_j^m C_j^m + \pi_4 D_j^m D_j^f + X_j \pi_x + \xi_j^f \tag{2.B0.2}$$

$$S_j^m = \psi_0 + \psi_1 D_j^f C_j^f + \psi_2 D_j^f C_j^m + \psi_3 D_j^m C_j^m + \psi_4 D_j^m C_j^f + X_j \psi_x + \xi_j^m \tag{2.B0.3}$$

where ξ_j^f, ξ_j^m are prediction errors. Here π_1 and $\pi_2$represent the increment in the probability of an actual female SGP if the vacancy (π_1 for the stereotypical, and π_2 for the non-stereotypical vacancies) relative to the omitted category of a prediction of no SGP. π_3 and π_4 represent the increment in the probability of an actual female SGP if the vacancy has a predicted male stereotypical SGP (π_3 for the stereotypical, and π_4 for the non-stereotypical vacancy), again relative to the case where it is predicted to have no SGP. Thus we expect π_1 and π_2 to be positive, whereas $\pi_3$and π_4 will be negative. Similar reasoning suggests that ψ_3 and ψ_4 will be positive and $\psi_1$and ψ_2 will be negative.

Combining equation (2.B0.1) with (2.B0.2) and (2.B0.3) shows that the difference-of-differences coefficients in (1) are:

$$\lambda_1 = \theta_1\pi_1 + \theta_4\psi_1 \tag{2.B0.4}$$

$$\lambda_2 = \theta_2\pi_2 + \theta_3\psi_2 \tag{2.B0.5}$$

$$\lambda_3 = \theta_3\psi_3 + \theta_2\pi_3 \tag{2.B0.6}$$

$$\lambda_4 = \theta_4\psi_4 + \theta_1\pi_4 \tag{2.B0.7}$$

Notice that if we ignore ψ_1 and ψ_2,then λ_1 is an attenuated version of θ_1 and λ_2 is an attenuated version of ψ_2, where the attenuation factors reflect the fractions of predicted vacancies with female stereotypical (π_1) or non-stereotypical (π_2) preferences (conditional on the $X's$). Also, if we ignore π_3 and π_4,then λ_3 is an attenuated version of θ_3 and λ_4 is an attenuated version of θ_4, where the attenuation factors reflect the fractions of predicted vacancies with male stereotypical (ψ_3) or non-stereotypical (ψ_3) preferences that actually have these preferences (conditional on the $X's$). More generally we would expect ψ_1, ψ_2 π_3, and π_4 to be small in magnitude (though in each case negative), so the intuition of the benchmark case remains true.

Columns 2 and 4 of Table 2.4 present estimates of equations (2.B0.3) and (2.B0.3) using the observed SGP's in the 1999-2003 period and predictions from our first (leave-out-mean based) classification model. We see that $\pi_1 = 0.171$, $\pi_2 = 0.230$ and, while $\psi_1 = -0.012$ and $\psi_2 = -0.082$. Thus, controlling for industry, occupation, and the firm's lagged gender composition, having a stereotypical female predicted SGP raises the probability of an SGP of that gender by 17 percentage points relative to a vacancy that is predicted to have no SGP, while having a male non-stereotypical SGP of the opposite gender lowers the probability by 1 percentage point relative to the no-SGP base.

2.C PREDICTION QUALITY

Figure 2.C1: Prediction Quality

(a) Regression: With LOO

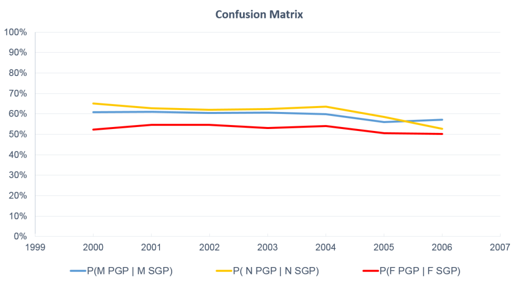

(b) Regression: Without LOO

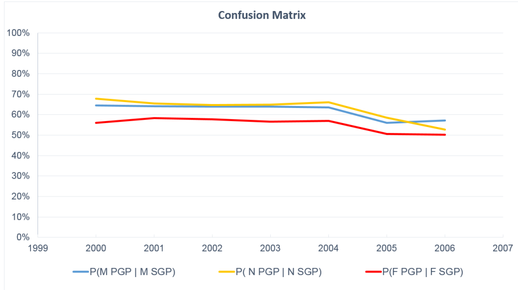

(c) Random Forest: With KFOLD

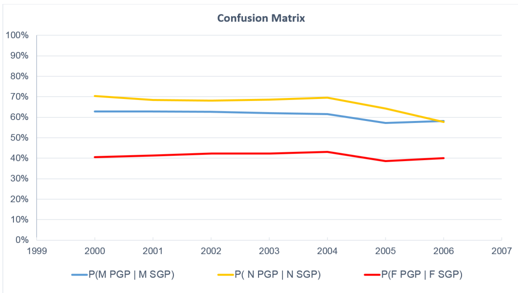

(d) Random Forest: Without KFOLD

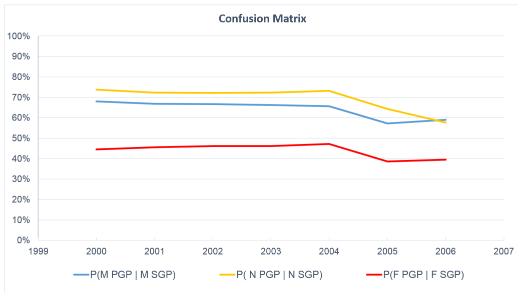

Notes: This figure shows the share of Stated Gender Preferences (SGP) that are correctly predicted by our Predicted Gender Preference (PGP), by year. The classification period is 2000-2004. Classification in panels (a) and (b) is performed through a fully saturated regression model, in panel (c) and (d) through the random forest model. In panel (b) we correct estimates through the leaveone-out mean mthod, in panel (c) we implement a rotating K-fold 1/4 vs 3/4 alghorithm.

Source: AMS-ASSD data, own calculations.

Table 2.C1: Predicted vs Observed SGP

Year	Observed	Prediction			
		Male PGP	**No PGP**	**Female PGP**	**Total**
2000	**Male SGP**	8,989	5,504	266	14,759
	No SGP	3,592	13,636	3,734	20,962
	Female SGP	237	5,328	6,130	11,695
	Total	12,818	24,468	10,130	47,416
2001	**Male SGP**	6,439	3,933	187	10,559
	No SGP	3,568	12,639	3,905	20,112
	Female SGP	127	3,942	4,896	8,965
	Total	10,134	20,514	8,988	39,636
2002	**Male SGP**	6,089	3,781	189	10,059
	No SGP	4,518	14,672	4,487	23,677
	Female SGP	148	3,833	4,782	8,763
	Total	10,755	22,286	9,458	42,499
2003	**Male SGP**	6,200	3,847	186	10,233
	No SGP	4,701	15,066	4,379	24,146
	Female SGP	147	4,279	5,010	9,436
	Total	11,048	23,192	9,575	43,815
2004	**Male SGP**	5,927	3,728	246	9,901
	No SGP	4,696	16,134	4,563	25,393
	Female SGP	139	3,951	4,818	8,908
	Total	10,762	23,813	9,627	44,202
2005	**Male SGP**	4,753	3,563	177	8,493
	No SGP	7,184	18,780	6,167	32,131
	Female SGP	170	3,742	4,011	7,923
	Total	12,107	26,085	10,355	48,547
2006	**Male SGP**	1,181	830	54	2,065
	No SGP	12,881	24,707	9,283	46,871
	Female SGP	56	736	798	1,590
	Total	14,118	26,273	10,135	50,526
2007	**Male SGP**	313	250	9	572
	No SGP	11,877	26,094	10,394	48,365
	Female SGP	9	233	230	472
	Total	12,199	26,577	10,633	49,409
2008	**Male SGP**	22	33	1	56
	No SGP	11,382	26,894	10,345	48,621
	Female SGP	7	45	32	84
	Total	11,411	26,972	10,378	48,761
2009	**Male SGP**	34	14	1	49
	No SGP	9,874	22,679	9,345	41,898
	Female SGP		42	30	72
	Total	9,908	22,735	9,376	42,019
2010	**Male SGP**	31	37		68
	No SGP	10,755	23,632	9,247	43,634
	Female SGP	2	27	18	47
	Total	10,788	23,696	9,265	43,749

2.D Additional Results

Figure 2.D1: Female Hiring - Weighted Regression Results

A. Leave-Out-Mean

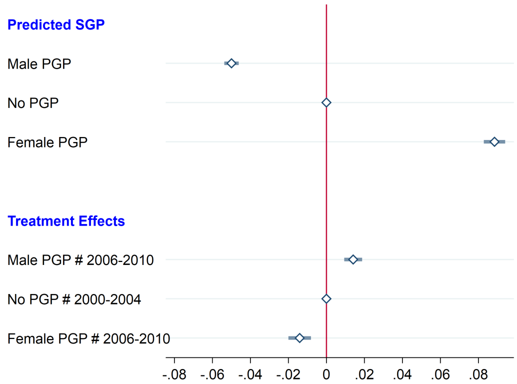

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females. The regressions follow the model in equation 2.3.1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects. Observations are weighted using representative sample regression weights. Classification uses regression.

Source: AMS-ASSD data, own calculations.

Figure 2.D2: Effects Vacancy Filling (controlling for Hire)

A. Baseline Model

B. Stereotypical vs Non-Stereotypical Vacancies

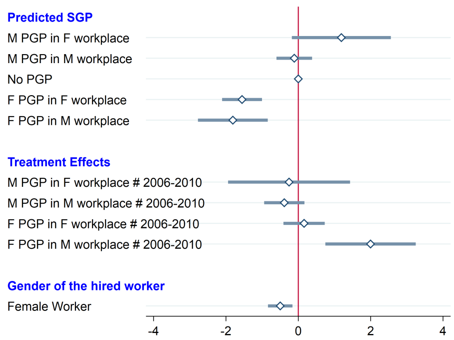

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on vacancy filling times. In panel A the the model follows the baseline specification in equation 2.3.1. In panel B the model follows the specification in equation 2.5.1, distinguishing between Stereotypical (M PGP in M Workplace and F PGP in F Workplace) and Non-Stereotypical (M PGP in F Workplace and F PGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include previous wage, gender, age, as well as industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

Figure 2.D3: Effects on Wages and Job Duration: Composition Effects (NOT controlling for Hire)

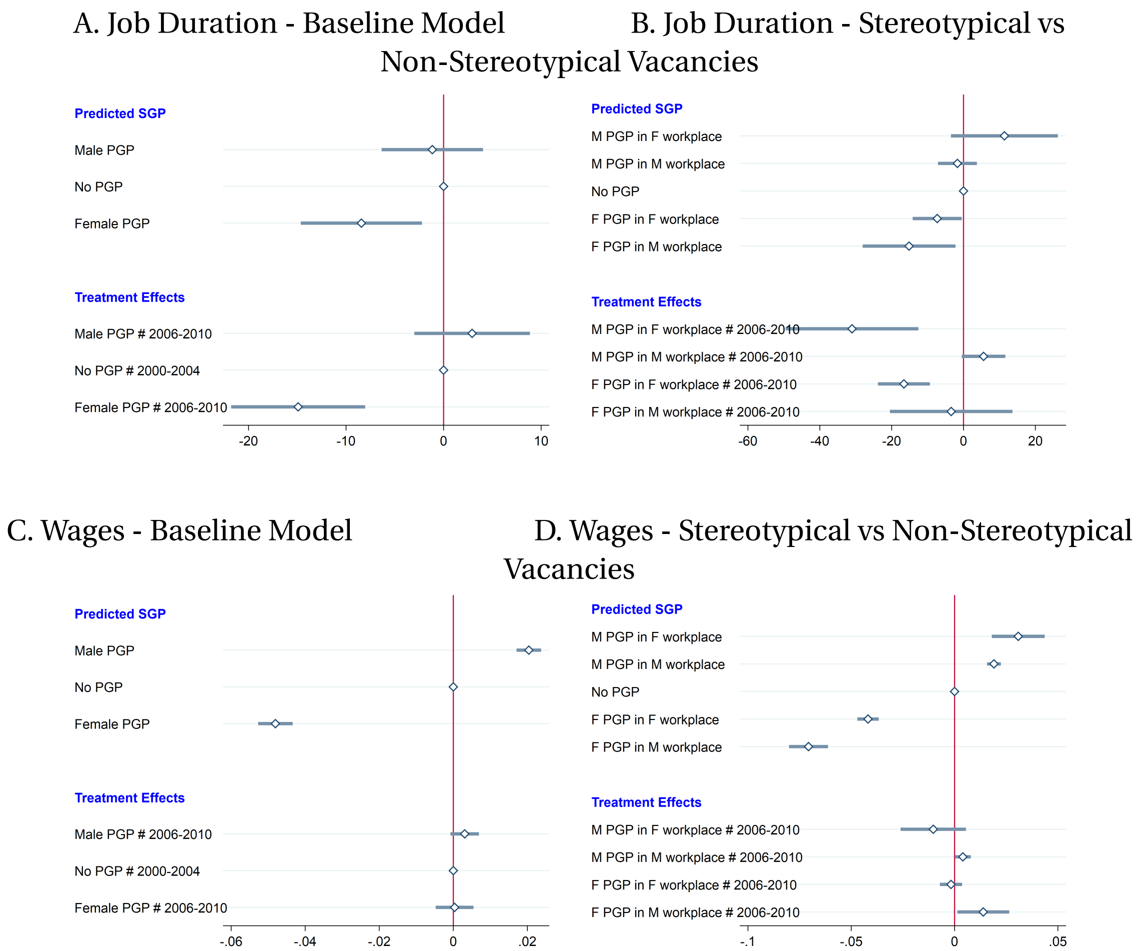

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on completed job durations (Panels A, B) and wages on the newly filled job (Panels C, D). The models follows the specification in equation 2.5.1, distinguishing between Stereotypical (M PGP in M Workplace and F PGP in F Workplace) and Non-Stereotypical (M PGP in F Workplace and F PGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

Figure 2.D4: Female Hiring (Difference in Differences Results)

Baseline + Industry-X-Year FE

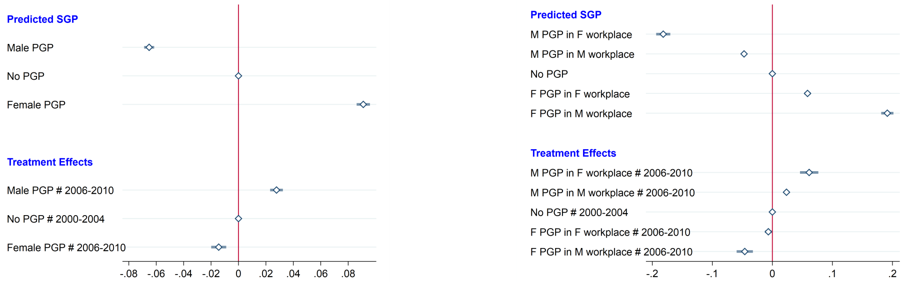

Baseline + Occupation-X-Year FE

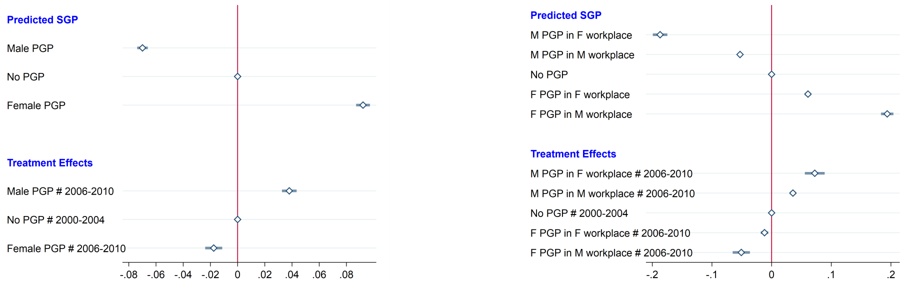

Baseline + Industry-X-Year FE + Occupation-X-Year FE

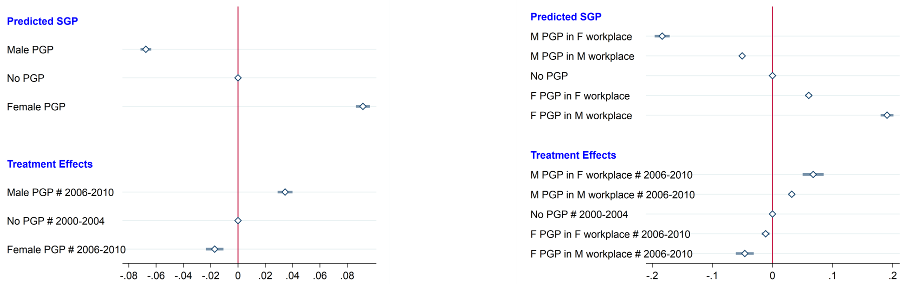

Baseline + Industry×Year Effects +
Occupation×Year Effects + FemaleShare×Year Effects

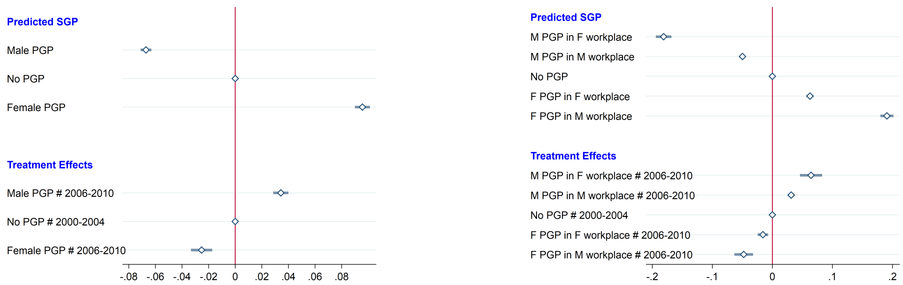

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on female hiring controlling for time variant industry, occupation and female share FE. The models follows the specification in equation2.3.1 in the left column and 2.5.1 in the right column. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped.

Source: AMS-ASSD data, own calculations.

Figure 2.D5: Effects on Job Duration (Event History Results)

Note: This figure reports the regression coefficients capturing the effect of eliminating stated gender preferences on completed job durations. Coefficients of the interaction term between year and indicators for job ads classified as advertising for women ("Female PGP"), are reported. Dotted lines show the 95% confidence intervals. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

*Source:*AMS-ASSD data, own calculations. calculations

Figure 2.D6: Effects on Workforce Diversity ($e_j \equiv H_j(1-C_j^f)+(1-H_j)C_j^f$)

A. Baseline Model B. Stereotypical vs Non-Stereotypical Vacancies

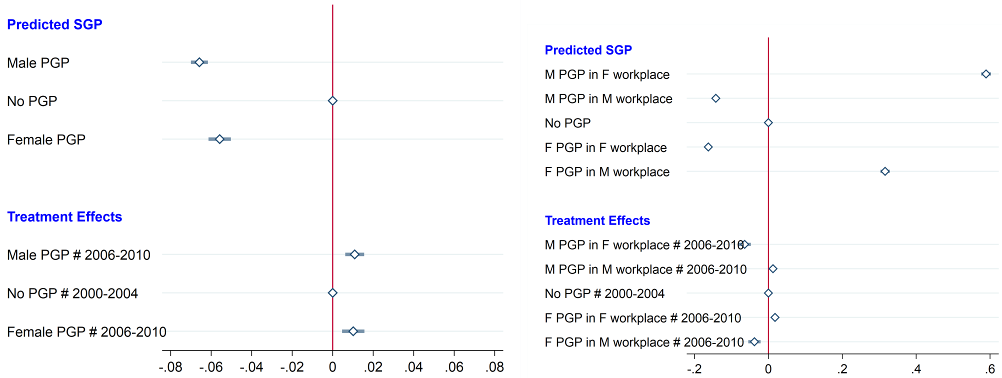

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on workplace composition using the discrete indicator. In panel A the model follows the baseline specification in equation 2.3.1. In panel B the model regression follows the specification in equation 2.5.1, distinguishing between Stereotypical (Predicted M SGP in M Workplace and F SGP in F Workplace) and Non-Stereotypical (Predicted M SGP in F Workplace and F SGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

2.E Tables main results

Table 2.E1: Effect of Eliminating Stated Gender Preferences on Female Hiring

	OLS Estimation				
Dependent Variable:	*Female Hiring*			*Workplace Diversity Index*	
	(1)	(2)	(3)	(4)	(5)
No PGP	omitted	omitted	omitted	omitted	omitted
Male PGP	-0.065***	-0.050***		-0.042***	
	(0.002)	(0.002)		(0.001)	
Female PGP	0.089***	0.089***		-0.034***	
	(0.002)	(0.002)		(0.002)	
Male PGP × 2006-2010	0.025***	0.013***		0.007***	
	(0.002)	(0.002)		(0.001)	
Female PGP × 2006-2010	-0.010***	-0.016***		0.006***	
	(0.002)	(0.003)		(0.002)	
M PGP in F workplace			-0.187***		0.211***
[Non-stereotypical Male PGP]			(0.006)		(0.004)
M PGP in M workplace			-0.047***		-0.070***
[Stereotypical Male PGP]			(0.002)		(0.001)
F PGP in F workplace			0.057***		-0.090***
[Stereotypical Female PGP]			(0.003)		(0.002)
F PGP in M workplace			0.193***		0.157***
[Non-stereotypical Female PGP]			(0.005)		(0.003)
M PGP in F workplace × 2006-2010			0.068***		-0.023***
[Non-stereotypical Male PGP × 2006-2010]			(0.008)		(0.005)
M PGP in M workplace × 2006-2010			0.023***		0.008***
[Stereotypical Male PGP × 2006-2010]			(0.002)		(0.001)
F PGP in F workplace × 2006-2010			-0.004		0.011***
[Stereotypical Female PGP × 2006-2010]			(0.003)		(0.002)
F PGP in M workplace × 2006-2010			-0.049***		-0.024***
[Non-stereotypical Female PGP × 2006-2010]			(0.007)		(0.004)
Observations	452029	452029	452029	452029	452029

Notes: OLS estimation results capturing the effect of eliminating stated gender preferences on the hiring of females (columns 1-3), and on the changes in workplace composition (columns 4-5). Regressions in columns 1,2 and 4 follow the specification in equation 2.3.1. Regressions in columns 3 and 5 follow the specification in equation 2.5.1. Classification in column 2 uses random forest. Industry, occupation, lagged female share in the workplace and year fixed effects are included as controls. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2004 are dropped from the regression. Beta coefficients reported and robust standard errors in parentheses. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

Source: AMS-ASSD data, own calculations.

Table 2.E2: Effect of Eliminating Stated Gender Preferences on Other Outcomes

	OLS Estimation					
Dependent Variable:	*Vacancy Filling*		*Daily Wage*		*Job Duration*	
	(1)	(2)	(3)	(4)	(5)	(6)
No PGP	omitted	omitted	omitted	omitted	omitted	omitted
Male PGP	-0.180		0.020***		-1.153	
	(0.198)		(0.002)		(2.653)	
Female PGP	-1.409***		-0.048***		-8.429***	
	(0.210)		(0.002)		(3.169)	
Male PGP × 2006-2010	0.024		0.003		2.921	
	(0.228)		(0.002)		(3.024)	
Female PGP × 2006-2010	0.538**		0.000		-14.915***	
	(0.232)		(0.003)		(3.505)	
M PGP in F workplace		0.686		0.031***		11.378
[Non-stereotypical Male PGP]		(0.546)		(0.007)		(7.595)
M PGP in M workplace		-0.292		0.019***		-1.693
[Stereotypical Male PGP]		(0.207)		(0.002)		(2.762)
F PGP in F workplace		-1.396***		-0.042***		-7.318**
[Stereotypical Female PGP]		(0.233)		(0.003)		(3.487)
F PGP in M workplace		-1.839***		-0.071***		-15.152**
[Non-stereotypical Female PGP]		(0.414)		(0.005)		(6.592)
M PGP in F workplace × 2006-2010		0.091		-0.010		-31.024***
[Non-stereotypical Male PGP × 2006-2010]		(0.677)		(0.008)		(9.427)
M PGP in M workplace × 2006-2010		0.014		0.004**		5.581*
[Stereotypical Male PGP × 2006-2010]		(0.234)		(0.002)		(3.099)
F PGP in F workplace × 2006-2010		0.313		-0.002		-16.587***
[Stereotypical Female PGP × 2006-2010]		(0.244)		(0.003)		(3.690)
F PGP in M workplace × 2006-2010		2.088***		0.014**		-3.429
[Non-stereotypical Female PGP × 2006-2010]		(0.541)		(0.006)		(8.701)
Observations	452029	452029	452027	452027	452029	452029

Notes: OLS estimation results capturing the effect of eliminating stated gender preferences on vacancy filling time (columns 1-2), on the wage of the newly filled job (columns 3-4) and on completed duration of the newly filled job (columns 5-6). Regressions in columns 1, 3 and 5 follow the specification of equation 2.3.1. Regressions in columns 2, 4 and 6 follow the specification in equation 2.5.1. Industry, occupation, lagged female share in the workplace and year fixed effects are included as controls. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2004 are dropped from the regression. Previous wage available only for 68 percent of the sample. Beta coefficients reported and robust standard errors in parentheses. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

Source: AMS-ASSD data, own calculations.

Table 2.E3: Effect of Eliminating Stated Gender Preferences on Other Outcomes, with controls for the Characteristics of the Hired Worker

	OLS Estimation					
Dependent Variable:	*Vacancy Filling*		*Daily Wage*		*Job Duration*	
	(1)	(2)	(3)	(4)	(5)	(6)
No PGP	omitted	omitted	omitted	omitted	omitted	omitted
Male PGP	0.023		0.006^{***}		6.082^{*}	
	(0.242)		(0.002)		(3.462)	
Female PGP	-1.515^{***}		-0.030^{***}		-9.899^{**}	
	(0.254)		(0.003)		(4.024)	
Male PGP × 2006-2010	-0.373		0.005^{**}		0.855	
	(0.277)		(0.002)		(3.931)	
Female PGP × 2006-2010	0.391		0.000		-18.109^{***}	
	(0.277)		(0.003)		(4.396)	
M PGP in F workplace		1.190^{*}		-0.021^{**}		28.680^{***}
[Non-stereotypical Male PGP]		(0.698)		(0.009)		(10.542)
M PGP in M workplace		-0.111		0.008^{***}		4.456
[Stereotypical Male PGP]		(0.251)		(0.002)		(3.583)
F PGP in F workplace		-1.556^{***}		-0.031^{***}		-6.297
[Stereotypical Female PGP]		(0.281)		(0.003)		(4.433)
F PGP in M workplace		-1.809^{***}		-0.026^{***}		-25.454^{***}
[Non-stereotypical Female PGP]		(0.492)		(0.005)		(8.261)
M PGP in F workplace × 2006-2010		-0.255		0.023^{**}		-39.175^{***}
[Non-stereotypical Male PGP × 2006-2010]		(0.860)		(0.011)		(13.037)
M PGP in M workplace × 2006-2010		-0.388		0.004^{*}		3.598
[Stereotypical Male PGP × 2006-2010]		(0.284)		(0.002)		(4.016)
F PGP in F workplace × 2006-2010		0.160		-0.000		-20.731^{***}
[Stereotypical Female PGP × 2006-2010]		(0.291)		(0.003)		(4.628)
F PGP in M workplace × 2006-2010		1.996^{***}		0.002		-0.907
[Non-stereotypical Female PGP × 2006-2010]		(0.638)		(0.007)		(10.721)
Female Worker	-0.507^{***}	-0.500^{***}	-0.179^{***}	-0.179^{***}	45.937^{***}	46.199^{***}
	(0.172)	(0.173)	(0.002)	(0.002)	(2.426)	(2.428)
Observations	308075	308075	308075	308075	308075	308075

Notes: OLS estimation results capturing the effect of eliminating stated gender preferences on vacancy filling time (columns 1-2), on the wage for the newly created job (columns 3-4), and on the duration of the newly created job (columns 5-6). Regressions in columns 1, 3 and 5 follow the specification of equation 2.3.1. Regressions in columns 2, 4 and 6 follow the specification of equation 2.5.1. Previous gender, wage and age of the hire, as well as industry, occupation, lagged female share in the workplace and year fixed effects are included as controls. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2004 are dropped from the regression. Previous wage available only for 68 percent of the sample. Beta coefficients reported and robust standard errors in parentheses. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

Source: AMS-ASSD data, own calculations.

CHAPTER 3

WHO BENEFITS FROM SUPPORT? THE HETEROGENEOUS EFFECTS OF SUPPORTERS ON ATHLETES' PERFORMANCE BY SKIN COLOR

3.1 INTRODUCTION

Racism at the workplace can have an impact on workers' wellbeing, which in turn affect their performance.[1] Productivity at work is often altered by the influence an individual receives from others (Kinlaw, 1999; Albrecht et al., 2014). Most companies condemn racist behaviors and pledge their commitment toward corporate social responsibility and social justice.[2] Despite the relevance of the phenomenon, research on the effect of racism on worker performance is limited. Racist episodes have been studied in some working environments, such as hospitals (Shields and Price, 2002) and the army (Antecol and Cobb-Clark, 2009). However, it is challenging to provide robust empirical evidence because of the difficulty to find an exogenous source of variation in exposure to racism.

[1] Oswald et al. (2015) show a link between happiness and productivity at work. Pagán-Castaño et al. (2020) provide an extensive literature review on the factors influencing employee wellbeing. Job satisfaction has been shown to be a good predictor of future stock market performance (Edmans, 2012), and increases value added (Bockerman and Ilmakunnas, 2012).

[2] Many private and public institutions fight against racism and try to increase awareness of the phenomenon among people. Several organizations seek to help workers and firms to detect racism harassment at the workplace and contribute to a more inclusive company culture. Examples are "Great Place to Work" in the US, "Pearn Kandola" in the UK and "Embrace Difference" in Europe.

The COVID-19 pandemic generated an important discontinuity in the way workers interact among them and with the external environment. Due to the COVID-19 restrictions, the main European leagues soccer games were played without fans in stadiums, which affects athletes behavior. Soccer has been shown to suit well for studying moral support as well as psychological, and social pressure generated from fans on athletes.[3] In soccer, racism is a widespread phenomenon and manifests mainly through discriminatory behaviors of supporter against non-white players.[4] The empty-seats shock reduces the positive effect of support from fans, but at the same time takes away the negative effect of racism on discriminated players. The availability of detailed data on soccer player performance and the prevalence of racism at stadiums make soccer a perfect environment to study the potential effects of racial offenses on individual performance at the workplace.

In this paper, I provide evidence that supporters have heterogeneous effects on soccer players performance by skin color in the highest Italian soccer league, Serie A. I exploit the sudden and unexpected change in the presence of supporters due to the COVID-19 restrictions. Soccer leagues discontinued the game schedule in March 2020 and restarted a few months later with one major change: the so called “closed stadium" rule, which forbids all individuals that are not directly involved in the game from entering stadiums. I generate and apply a skin color recognition algorithm to athlete pictures and classify more than 500 players into white and non-white categories. Using an objective performance measure from fantasy-sports competition, I compile individual performance scores in every game of the 2019/2020 season. I compare how each player fared in games played when supporters are allowed in the stadium against performance in empty stadiums.[5]

I find that performance of non-white players, relative to white players, increases

[3]See Colella et al. (2021); Gómez and Pollard (2011); Liardi and Carron (2011); Apesteguia and Palacios-Huerta (2010); Feri et al. (2013); Garicano et al. (2005) for pre-COVID-19 papers and Cross and Uhrig (2020); Dilger and Vischer (2020); Ferraresi and Gucciardi (2020); Fischer and Haucap (2020); Reade et al. (2020a); Sors et al. (2020); Bryson et al. (2021); Cueva (2020); Endrich and Gesche (2020); Scoppa (2021); Reade et al. (2020b) for research exploiting the “empty seats" discontinuity due to the COVID-19 pandemic.

[4]Thanks to new technology and an increasing awareness of the phenomenon, most of the racist episodes around soccer at an highly professional level are recorded. All the most important European leagues show an increasing pattern in racist episodes.

[5]Technically, the strategy is identical to a difference-in-differences methodology. However, in this setting there is no real control group as both “types" of players are affected by the absence of supporters, but in an heterogenous way.

by 1.5% on average, when fans are absent. The effects are similar in home and away games, and between players playing in top clubs and in other players. Defenders and midfielders, as well as less-skilled players, suffer more than others. Results are robust to the inclusion of several controls, such as turn, team and player fixed effects, game characteristics indicators and to potential confounders, such as player nationality. I also implement a placebo exercise using only 240 games played in front of fans: I generate a placebo ban, taking place after the first 120 games and replicate the main analysis. The placebo-ban has no significant effects on player performance score.

To the best of my knowledge, this is the first paper providing empirical evidence on the causal effect of supporters on individual performance by player skin color in a highly competitive environment.

This paper contributes to several strands of the literature. First, it relates to the literature connecting racism, wellbeing, and performance at work. Shields and Price (2002) document how nurses suffering racial harassment from patients report lower job satisfaction and are more intent on leaving their job, while Antecol and Cobb-Clark (2009) show that offensive racial behaviors toward military personnel heighten their intentions to leave the military profession. More recently, Stoermer et al. (2019) finds that black workers that experience racial harassment in the South African labor market have a lower job satisfaction and display a lower productivity. Turning to a firm level analysis, Corritore et al. (2020) show how workplaces with greater intrapersonal cultural heterogeneity between employees have greater expectations for future growth and innovation.[6] Given the difficulty to find an exogenous variation in racial harassment, research in this field is limited to descriptive evidence or survey studies. This paper contributes to this literature by producing new insights on the effect of external support, or pressure, on individual performance in a competitive working environment, by exploiting a natural experiment.

Second, this paper contributes to the economic research strand that uses soccer data to investigate human behavior.[7] Apesteguia and Palacios-Huerta (2010) collect data on penalty shoot-outs to show that being the first-mover increases the probability of winning, as a consequence of the psychological pressure suffered by second shoot-

[6]A closely related strand of the literature explores racial discrimination among workers and its impact on wages (Brown, 1984; Card and Krueger, 1993; Hersch and Shinall, 2015).

[7]The book *"Beautiful game theory: How soccer can help economics"* by Palacios-Huerta (2014) provides a review of previous research in economics using soccer data.

ers. Dohmen (2008) provides additional evidence on the effect of pressure on penalty score probability. Analyzing an exogenous shock in the presence of away team supporters in the Argentinean soccer league, Colella et al. (2021) identify an overall positive average effect of supporters on team performance. Previous research shows how home advantage, which consists in the greater probability of winning a game when playing at the local stadium, is present in team sports (Gómez and Pollard, 2011; Liardi and Carron, 2011; Carron et al., 2005; Pettersson-Lidbom and Priks, 2010; Pollard, 2006) as well as in individual competitions (Koning, 2011). It has been also shown that, in addition to directly affecting players, supporters exert social pressure on referees, biasing their decisions (Garicano et al., 2005; Dohmen and Sauermann, 2016). Recently, several studies have exploited the COVID-19 pandemic to provide additional evidence on home advantage (Cross and Uhrig, 2020; Dilger and Vischer, 2020; Ferraresi and Gucciardi, 2020; Fischer and Haucap, 2020; Reade et al., 2020a; Sors et al., 2020) and referee bias (Bryson et al., 2021; Cueva, 2020; Endrich and Gesche, 2020; Scoppa, 2021; Reade et al., 2020b;).[8] The literature on home advantage and the importance of supporters for performance looks mainly at team level effects, analyzing the total number of goals or cards given to each team as outcomes. This paper adds to this literature by providing evidence on the heterogeneous effect of supporters on individual player performance, skin color.

Third, this paper contributes to the literature on racism among supporters. Previous research in this field qualitatively documents the presence of racism among soccer fans by describing racist traits of supporters (Arnold and Veth, 2018; Hylton, 2020), or assessing their perception of xenophobic behavior of fans and media from the player prospective (Garland and Rowe, 2001). The book “European Football in Black and White” by Kassimeris (2007) describes the link between racism in soccer arenas and politics, and shows how the escalating violence, especially against ethnic and religious minorities, constitutes a threat to the emerging multicultural nature of European soccer. While these studies point toward a substantial presence of racism among fans, there is no evidence that supporters can affect player performance through racist behaviors. This paper fills this gap by showing empirically the negative causal effect of

[8]An article featured in The Economist (2020) under the headline “Home Comforts" (July 2020) provides a powerful graphical description of the changes in home advantage for the major European leagues.

supporters on the performance of non-white players.

The most closely related research to this paper is a working paper by Caselli et al. (2021). They study how the absence of supporters during soccer games impacts the performance of African Serie A players.[9] Their approach presents two main limitations compared to this study. First, they analyze the heterogeneous impact of supporters by player continent of origin. As result, their treatment group includes all white african players that might not be targeted by the crowd and, importantly, it excludes all non-white players with a European passport who are also victims of racism abuse.[10] Second, because of their citizenship based classification, their main result is based on a treatment group accounting for only 6.7% of the sample. I overcome these limitations by implementing a rigorous and objective classification of players following their skin color, resulting in a treatment group that accounts for 15.4% of the sample.

3.2 Background

Racism is present in many sports. In soccer, it manifests mainly through discriminatory behaviors and verbal abuse against non-white players. Most European soccer leagues have witnessed racist incidents in stadiums in the past, and the phenomenon is growing.[11] In recent years, awareness of the problems associated with racism has increased tremendously. Consequently, sports federations have moved jointly to fight racism: FIFA president Gianni Infantino has urged for "harsh sanctions" against clubs, while UEFA President Aleksander Ceferin has invited referees to stop games if a racist incident occurs. The idea of campaigning against racism in soccer is now rooted in many European countries. Since 2016 an ad of the "NO TO RACISM" campaign, which contains the most famous soccer players and managers as testimonials, has been shown before the start of any European league game.

The greatest portion of racist incidents in soccer concerns episodes of racial dis-

[9] In their paper, they refer to a The Economist (2021) article, that covered the findings of a previous version of this study, to highlight media interest in the topic and to motivate their research.

[10] Recent examples are racist abuse suffered by the English national team players Marcus Rashford, Jadon Sancho and Bukayo Sako after the Euro 2020 tournament final and by Italian player Mario Balotelli during a serie A game. A recent article on The Guardian (2021) stressed how the last Euro 2020 tournament, a competition played among national European teams, is a "barometer of diversity" due to the diverse composition of teams in terms of skin color. Romelu Lukaku, a Belgian serie A non-white player decleared in an interview:"*if it's going well, they call me a Belgian striker. If it's not, a Congolese descent*".

[11] Figures published by the UK government show a 50% increase in the number of incidents involving racist abuse during soccer games reported to the British police (The Guardian).

criminatory behaviors from in stadiums. In general, supporters are an integral part of the sport. They perceive themselves as part of the team and go to the stadium to boost players performance and to help the team to win games. They express their appreciation for or criticism against both own and opponent team players. Unfortunately, when a non-white player is targeted, fans behavior takes the form of racial harassment. Many professional players report having been victims of racist chants or verbal abuses by supporters during soccer games in recent years.[12]

Racist behaviors from supporters are unfortunately very common in Italy, where about 250 thousands supporters watch games live in stadiums every week. A report from the Italian soccer players association detected more than 600 racist incidents during the 2018/2019 season, 66% of them happened in stadiums. The victims are always non-white players, independent of their nationality. During a game of the 2019/2020 season, analyzed by this study, an Italian black soccer player, Mario Balotelli, stopped playing and threw the ball towards the supporter stands as a form of protest against the continuous insults he was receiving due to his skin color.

The current COVID-19 pandemic, has had a great impact on soccer. The most important soccer leagues discontinued the game schedule in March 2020 and restarted a few months later with one major change: the absence of supporters in stadiums. More than two thousand games, in the professional European leagues, were played with the closed stadium rule, which forbids all individuals that are not directly involved in the game, that is, players, managers, and staff from entering stadiums. The Italian Serie A resumed on June 20, 2020, until August 2, 2020. In total, 130 games, out of 380, were played in closed stadiums.

3.3 Data

3.3.1 Data Sources

The main source of data is *Fantacalcio.it* (last accessed: August 8, 2021), a platform for playing fantasy soccer in the Italian first league. Fantasy soccer is a fantasy sport game in which participants assemble imaginary or virtual teams of real professional

[12] Some examples are: Antonio Rudiger and Raheem Sterling in the English Premier League, Moussa Diaby in the German Bundesliga and Ousmane Dembele in the Spanish La Liga.

soccer players, and compete based on real player performance. After every Serie A game, Fantacalcio assigns a score between 0 and 10 to each player based on individual performance. The score is fully objective, as it is computed by the algorithm “Alvin482”, which examines several components of player behavior during the games including the position on the field, the share of successful passes, shoots, dribbles and tackles, as well as game specific characteristics such as the ranking of teams and the difficulty of the game. Each fantasy soccer team aims at getting the highest “Alvin482” total score. Fantacalcio also provides additional information on Serie A teams and players, such as age, weight, nationality, and position on the field, as well as results for every game played.

The second input concerns football players close-up photos downloaded from two websites: *Soccerwiki.org* and *Transfermarkt.com*. Soccerwiki is a free soccer-orientated wiki containing information on players, clubs, stadiums, managers, referees, leagues, and other data related to the world of soccer. Transfermarkt is a popular website collecting scores, results, and rankings of numerous leagues globally, as well as information on companies, players’ careers, and transfers. I downloaded more than 90% of the pictures from Soccerwiki. However, being a collaborative database where anyone can create and edit data, some of the photos are missing or “awaiting approval”. When one of these two events occurs, I downloaded the picture from the second source: Transfermarkt. At the end of this process I am able to obtain a picture for 557 Serie A players, accounting for 99.5% of the observations in the sample.

The last source of data is the website *Diretta.it*, a website collecting results, goals, and statistics of every game, as well as an indicator stating whether fans were allowed in the stadium when the game was played. I used this information to classify the sample into games with supporters and games without supporters.

3.3.2 Classification

After having downloaded every player photo, the algorithm classifies players into white and non-white.[13] To do that, I implement a three steps procedure. Figure (3.1) shows each step of the procedure for two players in the sample.

[13]The first category is supposed to contain all players from a light to an olive skin color, while the second includes players with a dark brown or black skin color.

Figure 3.1: Classification Examples

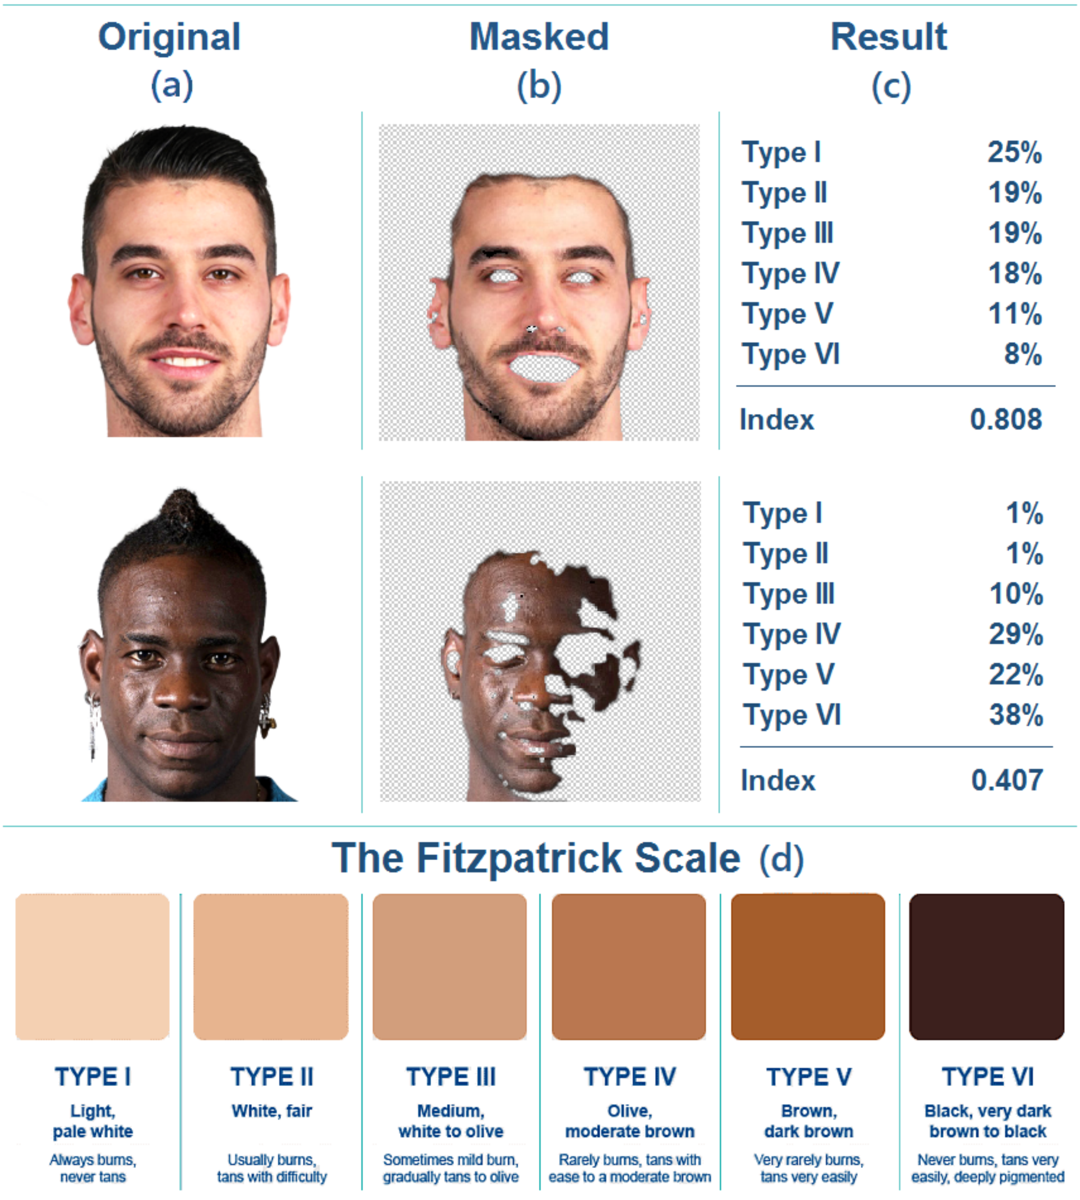

This Figure shows the three steps of the skin color classification algorithm for two Italian players in the sample. Panel (a) shows the picture downloaded from the website soccerwiki.com; panel (b) shows the figure after the mask, that isolates the skin, is applied; panel (c) shows the share of skin-pixels for each category of the Fitzpatrick scale and the skin-color-index generated; panel (d) plots the skin color categories in the Fitzpatrick scale - source: Fitzpatrick (1975).

First, I isolate all the pixels in a figure that identifies the player's skin. To this extent, I implement the novel *"skin detection using HSV & YCbCr color space"* method created by Dahmani et al. (2020).[14] Starting from an reed-green-blue (RGB) image, each pixel is converted into two color-composition scales: Hue Saturation Value (HSV) and Luma-Chroma-blue-Chroma-red (YCbCr). Then the value of each pixel is compared with standard values of a skin pixel and classified as skin and non-skin using a combination scheme.[15] This scheme exploits the weakness of each of the two skin detectors scales and improves the skin region segmentation by taking the best of each of the two classifiers. As shown in Figure (3.1), panel (b), on average between one third and half of the pixels are preserved.

Second, I classify each pixel in the figure in one of the 6 colors of the Fitzpatrick (1975) photo-typing scheme, presented in Figure (3.1) - panel (d). This is a numerical classification for human skin color commonly used in dermatology research. I use a color detection technique based on *OpenCV*, a library for image processing techniques implemented by Intel Corporation. The algorithm detects RGB coordinates for each pixel, computes the distance between the pixel coordinates and the coordinates of each of the 6 colors in the scheme, and selects the one having the smallest distance.[16] Panel (c) of Figure (3.1) shows the results of this procedure in terms of share of pixels for two players. As expected, the largest proportion of pixels for the first picture resides in the lighter categories, while the opposite happens in the case of the second picture.

The third and last step concerns the classification of players into one of the two categories: white and non-white. Given that many of the Italian or Mediterranean players have an "olive" skin color, I consider types I-IV for the white category and construct a skin color index that equals the share of pixels in these first four categories. As shown in Figure (3.1), the index for the first player is 0.8 while for the second it is 0.4. I classify as non-white all the players with an index value above 0.6 and as white all the others.[17] At

[14]They propose a novel skin segmentation method based on a zero-sum game theory model, which improves the detection efficiency of some usual skin descriptors. The algorithm used in this paper is freely accessible at the following address: https://github.com/CHEREF-Mehdi/SkinDetection. The website also contains an illustration on how the algorithm operates and several tests on its performance.

[15]The ranges are [$0 <= H <= 17$, $15 <= S <= 170$ and $0 <= V <= 255$] for the HSV scale and [$0 <= Y <= 255$, $135 <= Cr <= 180$ and $85 <= Cb <= 135$] for the YCbCr scale.

[16]The distance formula used is $[d_{pixel,i} = abs(Red_{pixel} - Red_{Fi}) + (Green_{pixel} - Green_{Fi}) + (Blue_{pixel} - Blue_{Fi})]$, where x_{pixel} refers to the color x coordinate of the pixel image and x_{Fi} refers to the color x coordinate of the i color in the Fitzpatrick scale.

[17]The distribution of the index presents a big jump after the cutoff of 0.6, with the majority of players almost equally distributed above it. I run a non-parametric regression of the grade on 9 indicators, one

the end of this procedure 86 of the 557 analyzed players in the sample were categorized as non-white; this accounts for 15.4% of the players and 14.2% of the observations.

3.3.3 Sample

Using the information provided by Fantacalcio for the entire 2019/2020 season of Serie A, and the skin color classification, I construct a panel database containing: (i) the score that each player received after each game, as well as (ii) baseline player characteristics, such as age, weight, height, nationality, team, and position on the field, and (iii) the assigned skin color. The database contains 20 teams, 557 players, and 380 games, for a total of 10,898 observations. The presence of non-white players is quite stable between home (14.3) and away games (14.0). The presence of non-white players is slightly greater in the top 6 teams (17.4%) than it is in others (14.4%). 73% of all the African players, 19% of the Latin Americans and only 9% of the European players are classified as non-white. Even though there is a substantial correlation between the skin color index and the continent of origin, only 32% of the non-white players in the sample have an African state as first nationality, while almost 50% of them are Europeans.

The goal of the analysis is to understand whether the lack of supporters has different effects on player performance according to skin color. I divide games in the database into two groups: games played in normal circumstances and games played in a closed stadium. Non-white players account for 13.8% of the observations in the first group and 14.8% of the observations in the second group. To avoid biased and volatile estimations due to a low number of observations, I include in the sample only players with more than 3 appearances in each of the two groups, reducing the sample to 9,495 observations. Figure (3.2) plots the average score that every player received by skin color and whether supporters were present or not in the stadium. If a game is played with supporters, white players get, on average, a score of 0.076 higher than non-white players, equals to an increase of 1.3%. On the other hand, if supporters are not present in the stadium during the game, non-white players get 0.03 points more than white players, indicating a potential heterogeneity in the effect of supporters on players performance related to skin color.

for each 0.1 increment of the skin color index, and find that the effect picks up at the cutoff of 0.6. These results are available upon request.

Figure 3.2: Average Individual Score

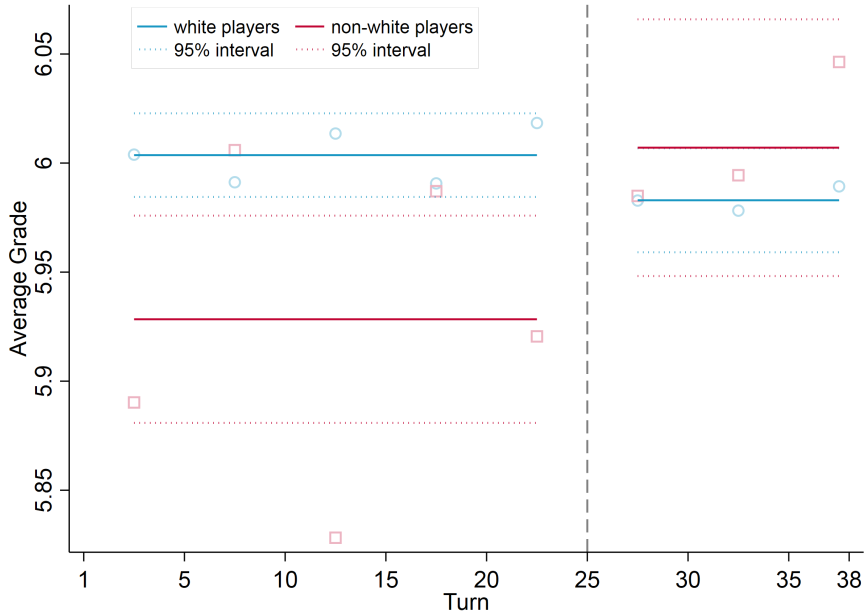

This Figure presents the average score received by white players (in blue) and non-white players (in red) in the 2019/2020 serie A season. The left side concerns the period in which games were played with supporters and the right one the period in which supporters were banned from stadium. The circles and the squares plot the averages by bins of 5 turns, 50 games. Source: Fantacalcio.it, own calculation.

3.4 Impact Estimates

3.4.1 Empirical Strategy

In order to causally identify the effect of (the lack of) supporters on the performance of white and non-white players, I estimate the following equation:

$$Y_{ikjt} = \alpha + \beta\, NW_i + \gamma\, NS_{jt} + \lambda\, NW_i \times NS_{jt} + \gamma\, X_{itkm} + \rho_i + \nu_t + \epsilon_{ikjt} \tag{3.0.1}$$

Where Y_{ikjt} refers to the score attributed to player i in team k during game j played for the Serie A turn t; NW_i is a dummy indicating whether player i is categorized as non-white and NS_{jt} indicates whether the game j was played without supporters in the stadium. Depending on the specification used, X_{itm} contains several controls, such as indicators for home games, team fixed effects, number of goals scored by team k in game j, a quadratic time control and the nationality of player i. ρ_i and ν_t are individual and time fixed effects, respectively.[18] The main coefficient of interest is the interaction term λ, which represents the average change in the performance score of a non-white player when there are no-supporters to when there are supporters in the stadium, relative to the same average change for white players.

Identification relies on two assumptions. The first is the absence of other factors that could affect the performance of non-white players, over that of white players, contemporaneously with the closed stadium policy. The second is a common trend assumption between the performance score of white and non-white players: I assume that if the closed stadium policy would have never been implemented, the difference in score between the two groups of players would have remained constant. Figure (3.2) doesn't show any particular trends in the two distributions. However, one might think that other unobserved individual characteristics of players, correlated with the skin color, might affect their performance and bias the results. I address potential concerns by running several robustness checks. In particular, in a double interaction model, I show that if the skin color is taken into account, the continent of origin does not significantly affect player performance. In addition, I run a placebo experiment on a subsample of games played with supporters in the stadium and do not find any significant result of the placebo ban implemented after 12 games, out of 24, in this hypothetical scenario.

3.4.2 Main Result: Effect on Player Performance

Table (3.1) shows the estimation results. In the most plain-vanilla specification, column (1), the non-white coefficient is always negative, indicating that non-white players receive, on average, a lower score with respect to white players. This difference can be due to a baseline disparity in skills and strengths between players in the two categories. The interaction term between the non-white and the no-supporters dummy is instead positive, suggesting that, without supporters, the average score of non-white players increases compared with the change that white players had. I consider a few different specifications. Column (2) controls for a quadratic time trend.[19] Then, I con-

[18]In all the specifications where a player fixed effect is included, the variable NW_i, as well as the player time invariant controls are omitted, while the variable NS_{jt} also appears in the specification including turn fixed effects because at the very beginning of the restrictions a few games of the same turn were played without supporters.

[19]By time, in this setting, I refer to turn, an indicator of the chronological order of the games.

trol for game characteristics, by adding the home game dummy and the number of goals scored by the team, in column (3), and team fixed effects, in column (4). To reduce potential bias from player time-invariant characteristics, I include player nationality fixed effects in column (6) and player fixed effects in column (7). Finally, I control for potential time non-parametric effects in column (8) by adding a turn fixed effect. In column 6 to 8, I also cluster standard errors at the player level, to obtain robust inference to potential unobserved correlation between different games played by the same player (Cameron and Miller, 2015).

Table 3.1: Effect of supporters on players performance

OLS Estimation								
Dependent Variable: *Player performance score*								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
No Supporters	-0.021	-0.049	-0.050*	-0.055*	-0.059**	-0.059*	-0.053*	0.053
	(0.015)	(0.032)	(0.030)	(0.030)	(0.030)	(0.032)	(0.032)	(0.068)
Non-white player	-0.072**	-0.072**	-0.089***	-0.100***	-0.102***	-0.102**		
	(0.028)	(0.028)	(0.026)	(0.027)	(0.034)	(0.052)		
No Supporters × Non-white player	0.096**	0.095**	0.093**	0.088**	0.087**	0.087*	0.090*	0.089*
	(0.044)	(0.044)	(0.040)	(0.040)	(0.040)	(0.047)	(0.047)	(0.047)
Turn Quadratic Trend		✓	✓	✓	✓	✓	✓	
Home Games Dummy			✓	✓	✓	✓	✓	✓
Team Goals			✓	✓	✓	✓	✓	✓
Team FE				✓	✓	✓	✓	✓
Player's Nationality FE					✓	✓		
Player FE							✓	✓
Turn FE								✓
Observations	8952	8952	8952	8952	8952	8952	8952	8952
Cluster SE - player						✓	✓	✓

OLS estimation of the effect of no supporters on player performance by skin color. No Supporters is a dummy taking value 1 for all games played in empty stadiums. Non-white player is a dummy taking value 1 for all players classified as non-white. Beta coefficients reported and robust standard errors in parentheses. Standard errors are clustered by player in Columns (6), (7) and (8). *** significant at 1%, ** significant at 5%, * significant at 10%.

Independently of the specification, the interaction term between the no-supporters and the non-white indicator is always positive, statistically significant, and stable in terms of magnitude, ranging between 0.087 and 0.096. In the most conservative specification, column (8), the coefficient is 0.089, indicating a 1.5% increment in the average score.[20] This result is consistent withe the hypothesis that non-white players suffer from the presence of supporters in the stadium. The COVID-19 pandemic may not only affect the presence of fans in stadiums, it could also have a direct effect on team performance in other ways. However, it is unlikely that other potential factors, other

[20]The pre-ban average score was 5.93 for non-white players and 6.01 for white players.

than racist behaviors from supporters, would impact differently the performance of players by skin color.

3.4.3 Robustness

In this section, I present the additional robustness checks. The first set of checks concerns sample selection. In the main analysis, I consider only players playing more than 3 games in each part of the season. In table (3.2), I replicate the most conservative specification, including time and player fixed effects, using all the players in the sample, column (1), and only players making at least 10 appearances per period, column (2). I then exclude the goalkeepers from the sample as they are all categorized as white, column (3). The interaction coefficient is always positive, statistically significant and quite stable in terms of magnitude, indicating that the result is not due to sample selection.

As a second check, I exclude the possibility that the result is due to potential unobserved characteristics of the players related to the origin country that are, at the same time, correlated with the impact of the closed stadium and the skin color. For this reason, I use the nationality to divide players into four subgroups according to their continent of origin.[21] In column (4), I replicate the main specification by including the interaction terms between the no-supporters indicator and each continent dummy, and the double interaction no-supporters, non-white and continent. While the main treatment effect remains positive and significant, none of the coefficients of the interactions with the continent indicators is significant, showing that, once the skin color is taken into account, the continent has no treatment effect on performance. In addition, the double interaction terms are not significant, demonstrating the absence of heterogeneous effects of the treatment by player continent. This evidence reinforces the theory that racist supporters target players because of their skin color, rather than their citizenship.

The last check deals with a placebo test to exclude that non-white players exhibit any difference in performance due to time rather than the closed stadium policy. I

[21] The sample contains a small amount of players with more than one nationality, in this case I consider only the first nationality mentioned as is intended to be the one players refer to as “first". As shown in the online appendix table A1, 76% of players are European, 16% are from Latin America, 6.7% are Africans and 1.3% Asiatic.

Table 3.2: Effect of supporters on player performance - Robustness Checks

	OLS Estimation				
Dependent Variable: *Player performance score*					
	Sample Selection			Continent	Placebo
	(1)	(2)	(3)	(4)	(5)
No Supporters ×Non-white player	0.077* (0.046)	0.090* (0.049)	0.081* (0.047)	0.124** (0.062)	
No Supporters×Africa				0.062 (0.057)	
No Supporters×America				0.037 (0.047)	
No Supporters×Asia				0.200 (0.129)	
No Supporters×Africa×Non-white player				-0.116 (0.117)	
No Supporters×America×Non-white player				-0.096 (0.117)	
No Supporters - Placebo×Non-white player					-0.009 (0.053)
Controls	✓	✓	✓	✓	✓
Player FE	✓	✓	✓	✓	✓
Turn FE	✓	✓	✓	✓	✓
Observations	9642	7751	8332	8332	5435
Cluster SE - player	✓	✓	✓	✓	✓
Sample	All	> 10 Games	No GK	Main	No Covid

OLS estimation of the effect of no supporters on player performance by skin color of the players. No Supporters is a dummy taking value 1 for all games played in empty stadiums. Non-white player is a dummy taking value 1 for all players classified as non-white. The sample in column (1) includes all players, in column (2) only players playing at least 10 games per period and in column (3) all players but goalkeepers. In column (5) only games played from turn 1 to turn 24 are included. The reference category for Continent is Europe (category with the highest number of observations). All players from Asia are classified as white so the interaction term with non-white is omitted. No Supporters - Placebo is a dummy taking value 1 for all games played from turn 13 to turn 24. Controls include Home games dummy, Team goals, Team FE, and Turn FE in all the specifications. It also includes No supporters dummy in columns (1)-(4) and No Supporters - Placebo in column (5). Beta coefficients reported and standard errors clustered by player in parentheses. *** significant at 1%, ** significant at 5%, * significant at 10%.

consider only games played from turn 1 to turn 24.[22] I create a placebo-ban variable taking value 1 only for games played after turn 13 and I replicate the main analysis. Column (5) displays the results. As expected the placebo-ban coefficient is null, providing supportive evidence that there was no differential trend between white and non-white players before the closed stadium policy.

3.4.4 Heterogeneous Effects

In this section, I investigate potential heterogeneous effects of the supporters on non-white players by characteristics of the game or the players. To this extent, I construct binary indicators for each characteristic and I replicate the main specification by including: (i) the indicator, (ii) the interaction term between indicator and the no-supporters dummy, (iii) the interaction term between indicator and the non-white dummy and (iv) the double interaction indicator, no-supporters, non-white. Table (3.3) displays the results.

To explore eventual heterogeneity of the effect with respect to the characteristics of the game, I consider indicators for home games and top teams. Previous research shows the existence of the home advantage and that it persists also when supporters are not in the stadium. Top teams is a dummy taking value 1 if the player plays in one of the top 6 teams.[23] As shown in columns (1-3), there is no statistically significant heterogeneous treatment with respect to game characteristics, signaling that non-white players suffer from supporters more than others independently of the difficulty of the game characteristics.

Finally, I investigate heterogeneous treatment effects by player role and player distribution. Results in columns (4) show that, in comparison with other players, strikers are not affected by the absence of supporters. Column (5) indicates that players in the third quartile of the pre-ban score distribution are significantly less affected than the one in the first quartile. In terms of magnitude, this completely offsets the baseline negative effect. While not significant, also coefficients for the second and fourth quartile are negative, signaling that the positive and statistically significant effect found in

[22] About half of the games in turns 25 and 26 were played without supporters, while all games form turn 27 on were played in closed stadiums.

[23] Top 6 teams are decided following popularity, number of supporters and recent years performance. They are: Juventus, Milan, Inter, Napoli, Roma and Lazio.

Table 3.3: Heterogeneous effects of supporters on player performance

OLS Estimation

Dependent Variable: *Player performance score*

	Game characteristics			Player's role	Player's strength
	(1)	(2)	(3)	(4)	(5)
No Supporters ×Non-white player	0.097*	0.113**	0.111*	0.157**	0.138**
	(0.055)	(0.049)	(0.059)	(0.064)	(0.060)
Home game×No Supporters×Non-white player	-0.014		0.004		
	(0.065)		(0.082)		
Top team×No Supporters×Non-white player		-0.079	-0.046		
		(0.116)	(0.131)		
Top team×Home game×No Supporters×Non-white player			-0.065		
			(0.115)		
Defender ×No Supporters×Non-white player				0.018	
				(0.092)	
Striker×No Supporters×Non-white player				-0.225**	
				(0.111)	
2nd Quartile×No Supporters×Non-white player					-0.071
					(0.086)
3rd Quartile×No Supporters×Non-white player					-0.198*
					(0.101)
4th Quartile×No Supporters×Non-white player					-0.060
					(0.127)
Controls	✓	✓	✓	✓	✓
Player FE	✓	✓	✓	✓	✓
Turn FE	✓	✓	✓	✓	✓
Observations	8952	8952	8952	8952	8944
Cluster SE - player	✓	✓	✓	✓	✓

OLS estimation of the effect of no supporters on players performance by skin color of the players. No Supporters is a dummy taking value 1 for all games played in empty stadiums. Non-white player is a dummy taking value 1 for all players classified as non-white. Top team is a dummy taking value 1 for all players playing for Juventus, Inter Milan, A.C. Milan, Napoli, Rome and Lazio. The reference category for role is midfielder (category with the highest number of observations) and for quartiles is the first quartile. All goalkeepers are classified as white so the interaction term with non-white is omitted. Controls include Home games dummy, Team goals, Team FE, and Turn FE in all the specifications. It also includes No supporters dummy and Home game × No supporters in columns (1) and (3), Top team × No supporters in columns (1) and (3), Home game × Top team in column (3), dummy for each role × No supporters in column (4), and dummy for each quartile × No supporters in column (5). Beta coefficients reported and standard errors clustered by player in parentheses. *** significant at 1%, ** significant at 5%, * significant at 10%.

the main analysis is coming from the bottom quartile of the player strength distribution.

3.5 Conclusion

In this paper, I investigate the effect of supporters in stadiums on the performance of soccer players by skin color, in the highest Italian soccer league, Serie A. I evaluate player performance using an objective score and classify players in white and non-white implementing an automated skin color recognition algorithm. Identification comes from an exceptional change in access to stadiums: due to the COVID-19 restrictions, Serie A 2019/2020 season interrupted in March 2020 and restarted in June 2020 allowing access to stadium only to players, managers and staff. As a result, one third of the games were played without fans. I find a significant increase by 1.5% in the performance of non-white players, relative to white players, when fans are not in the stadium. Given the absence of factors other than racist chants and verbal abuse from supporters that could impact differently players by skin color, this result suggest that racial discrimination faced by non-white players affects their performance. This study provides the first evidence on the heterogeneous effects of supporters on performance by player skin color. It complements the literature on moral support and performance and provides insights on the effect of racist behavior on performance, contributing to the labor economics literature connecting racist episodes, wellbeing and individual behavior at work.

APPENDIX

3.A DESCRIPTIVE STATISTICS

Table 3.A1: Descriptive Statistics

	Total	**By skin color**		
		White	**Non-white**	**Share**
	(1)	(2)	(3)	(4)
Players				
Number of Players	557	471	86	15.4%
By Continent of Origin (Passport)				
Africa	37	10	27	73.0%
America	89	72	17	19.1%
Europe	423	381	42	9.9%
Asia	8	8	0	0.0%
By Team				
Top 6 teams	172	142	30	17.4%
All other teams	389	333	56	14.4%
Observations - player X game				
Games with supporters	6,936	5,982	954	13.8%
Games without supporters	3,962	3,375	587	14.8%
Home Games	5,449	4,669	780	14.3%
Away Games	5,449	4,688	761	14.0%

This table reports the number of players in total, by continent of origin and by team and the number of observations in the sample for games played with supporters and with empty stadiums and for home and away games in column (1). Columns (2) and (3) reports the same statistics by the skin color of the Player and column (4) the share of non-white players or observations.

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.

_ **and David Autor**, "Skills, tasks and technologies: Implications for employment and earnings," in Orley Ashenfelter and David Card, eds., *Handbook of labor economics*, Vol. 4, Elsevier, 2011, pp. 1043–1171.

_ **and Pascual Restrepo**, "Artificial intelligence, automation and work," Technical Report, National Bureau of Economic Research 2018.

_ **and** _ , "The race between man and machine: Implications of technology for growth, factor shares, and employment," *American Economic Review*, 2018, *108* (6), 1488–1542.

Agan, Amanda and Sonja Starr, "Ban the box, criminal records, and racial discrimination: A field experiment," *The Quarterly Journal of Economics*, 2018, *133* (1), 191–235.

_ **, Bo Cowgill, and Laura Katherine Gee**, "Do workers comply with salary history bans? a survey on voluntary disclosure, adverse selection, and unraveling," in "AEA Papers and Proceedings," Vol. 110 2020, pp. 215–19.

Albrecht, Konstanze, Johannes Abeler, Bernd Weber, and Armin Falk, "The brain correlates of the effects of monetary and verbal rewards on intrinsic motivation," *Frontiers in Neuroscience*, 2014, *8*, 303.

Alvarez, Roberto and Ricardo A. Lopez, "Foreign Technology Acquisition and Changes in the Real Exchange Rate," *The World Economy*, 2015, *38* (4), 613–628.

Antecol, Heather and Deborah Cobb-Clark, "Racial harassment, job satisfaction, and intentions to remain in the military," *Journal of Population Economics*, 2009, *22* (3), 713–738.

Apesteguia, Jose and Ignacio Palacios-Huerta, "Psychological Pressure in Competitive Environments: Evidence From a Randomized Natural Experiment.," *American Economic Review*, 2010, *100* (5), 2548 –2564.

Arnold, Richard and Karl Manuel Veth, "Racism and Russian Football Supporters' Culture: A Case for Concern?," *Problems of Post-Communism*, 2018, *65* (2), 88–100.

Atalay, Enghin, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum, "The Evolving U.S. Occupational Structure," 2018, p. 95.

Auer, Raphael, Ariel Burstein, and Sarah M. Lein, "Exchange Rates and Prices: Evidence from the 2015 Swiss Franc Appreciation," *American Economic Review*, February 2021, *111* (2), 652–86.

_ , _ , **Katharina Erhardt, and Sarah M. Lein**, "Exports and Invoicing: Evidence from the 2015 Swiss Franc Appreciation," *AEA Papers and Proceedings*, May 2019, *109*, 533–38.

Autor, David H, "The'task approach'to labor markets: an overview," 2013.

_ **and David Dorn**, "The growth of low-skill service jobs and the polarization of the US labor market," *American Economic Review*, 2013, *103* (5), 1553–97.

_ , _ , **and Gordon H Hanson**, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

_ , _ , _ , **and Jae Song**, "Trade adjustment: Worker-level evidence," *The Quarterly Journal of Economics*, 2014, *129* (4), 1799–1860.

Autor, David H., Frank Levy, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration*," *The Quarterly Journal of Economics*, 11 2003, *118* (4), 1279–1333.

Azar, Jose, Ioana Elena Marinescu, Marshall Steinbaum, and Bledi Taska, "Concentration in US Labor Markets: Evidence from Online Vacancy Data," SSRN Scholarly Paper ID 3133344, Social Science Research Network, Rochester, NY August 2018.

Azmat, Ghazala and Barbara Petrongolo, "Gender and the labor market: What have we learned from field and lab experiments?," *Labour Economics*, 2014, *30* (C), 32–40.

Baert, Stijn, "Hiring Discrimination: An Overview of (Almost) All Correspondence Experiments Since 2005," GLO Discussion Paper Series 61, Global Labor Organization (GLO) 2017.

Ballance, Joshua, Robert Clifford, and Daniel Shoag, ""No more credit score": Employer credit check bans and signal substitution," *Labour Economics*, 2020, *63*, 101769.

Bastos, Paulo, Joana Silva, and Eric Verhoogen, "Export Destinations and Input Prices," *American Economic Review*, February 2018, *108* (2), 353–92.

Baurle, Gregor, Matthias Gubler, and Diego R. Kanzig, "International inflation spillovers - the role of different shocks," *Swiss National Bank Working Papers*, 2017, (7).

Bayard, Kimberly, Judith Hellerstein, David Neumark, and Kenneth Troske, "New evidence on sex segregation and sex differences in wages from matched employee-employer data," *Journal of labor Economics*, 2003, *21* (4), 887–922.

Becker, Sascha O, Karolina Ekholm, and Marc-Andreas Muendler, "Offshoring and the onshore composition of tasks and skills," *Journal of International Economics*, 2013, *90* (1), 91–106.

Beller, Andrea H, "The Effects of Title VII of the Civil Rights Act of 1964 on Women's Entry into Nontraditional Occupations: An Economic Analysis," *Law & Ineq.*, 1983, *1*, 73.

Bergmann, Barbara R, "Occupational segregation, wages and profits when employers discriminate by race or sex," *Eastern economic journal*, 1974, *1* (2), 103–110.

Biscourp, Pierre and Francis Kramarz, "Employment, skill structure and international trade: Firm-level evidence for France," *Journal of International Economics*, 2007, *72* (1), 22–51.

Blau, Peter Michael, *Inequality and heterogeneity: A primitive theory of social structure*, Vol. 7, Free Press New York, 1977.

Blinder, Alan S and Alan B Krueger, "Alternative measures of offshorability: a survey approach," *Journal of Labor Economics*, 2013, *31* (S1), S97–S128.

Bloom, Nicholas, Mirko Draca, and John Van Reenen, "Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity," *The review of economic studies*, 2016, *83* (1), 87–117.

Bockerman, Petri and Pekka Ilmakunnas, "The job satisfaction-productivity nexus: A study using matched survey and register data," *ILR Review*, 2012, *65* (2), 244–262.

Bonadio, Barthélémy, Andreas M Fischer, and Philip Sauré, "The Speed of Exchange Rate Pass-Through," *Journal of the European Economic Association*, 03 2019, *18* (1), 506–538.

Booth, Alison and Andrew Leigh, "Do employers discriminate by gender? A field experiment in female-dominated occupations," *Economics Letters*, May 2010, *107* (2), 236–238.

Bos, Marieke, Emily Breza, and Andres Liberman, "The labor market effects of credit market information," *The Review of Financial Studies*, 2018, *31* (6), 2005–2037.

Botosaru, Irene and Federico H Gutierrez, "Difference-in-differences when the treatment status is observed in only one period," *Journal of Applied Econometrics*, 2018, *33* (1), 73–90.

Brown, Charles, "Black-white earnings ratios since the Civil Rights Act of 1964: The importance of labor market dropouts," *The Quarterly Journal of Economics*, 1984, *99* (1), 31–44.

Bryson, Alex, Peter Dolton, J James Reade, Dominik Schreyer, and Carl Singleton, "Causal effects of an absent crowd on performances and refereeing decisions during Covid-19," *Economics Letters*, 2021, *198*, 109664.

Burstein, Ariel, Javier Cravino, and Jonathan Vogel, "Importing skill-biased technology," *American Economic Journal: Macroeconomics*, 2013, *5* (2), 32–71.

Bustos, Paula, "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *American Economic Review*, February 2011, *101* (1), 304–40.

Cahuc, Pierre, Stéphane Carcillo, Andreea Minea, and Marie-Anne Valfort, "When Correspondence Studies Fail to Detect Hiring Discrimination," CEPR Discussion Papers 14028, C.E.P.R. Discussion Papers September 2019.

Cameron, A Colin and Douglas L Miller, "A practitioner's guide to cluster-robust inference," *Journal of Human Resources*, 2015, *50* (2), 317–372.

Campa, Jose Manuel and Linda S Goldberg, "Employment versus wage adjustment and the US dollar," *Review of Economics and Statistics*, 2001, *83* (3), 477–489.

Card, David, "Using regional variation in wages to measure the effects of the federal minimum wage," *Ilr Review*, 1992, *46* (1), 22–37.

_ , "Who Set Your Wage?," Technical Report, National Bureau of Economic Research 2022.

_ , Ana Rute Cardoso, and Patrick Kline, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *The Quarterly Journal of Economics*, May 2016, *131* (2), 633–686.

_ and Alan B Krueger, "Trends in relative black-white earnings revisited," *The American Economic Review*, 1993, *83* (2), 85–91.

_ 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.

Carron, Albert V, Todd M Loughhead, and Steven R Bray, "The home advantage in sport competitions: Courneya and Carron's (1992) conceptual framework a decade later," *Journal of Sports Sciences*, 2005, *23* (4), 395–407.

Caselli, Mauro, Paolo Falco, and Gianpiero Mattera, "When the Stadium Goes Silent: How Crowds Affect the Performance of Discriminated Groups," Technical Report, Department of Economics and Management, 06, 2021.

Chen, Shiu-Sheng, "Exchange rate undervaluation and R&D activity," *Journal of International Money and Finance*, 2017, *72* (C), 148–160.

Colella, F, Patricio Dalton, and G Giusti, "All you Need is Love: The Effect of Moral Support on Performance (Revision of CentER DP 2018-026)," 2021.

Corritore, Matthew, Amir Goldberg, and Sameer B Srivastava, "Duality in diversity: How intrapersonal and interpersonal cultural heterogeneity relate to firm performance," *Administrative Science Quarterly*, 2020, *65* (2), 359–394.

Cortes, Patricia and Jessica Pan, "Occupation and Gender," IZA Discussion Papers 10672, Institute of Labor Economics (IZA) March 2017.

Cross, Jeffrey and Richard Uhrig, "Do Fans Impact Sports Outcomes? A COVID-19 Natural Experiment," *A COVID-19 Natural Experiment (October 5, 2020), Mimeo*, 2020.

Cueva, Carlos, "Animal Spirits in the Beautiful Game. Testing social pressure in professional football during the COVID-19 lockdown," *Mimeo*, 2020.

Dahmani, Djamila, Mehdi Cheref, and Slimane Larabi, "Zero-sum game theory model for segmenting skin regions," *Image and Vision Computing*, 2020, *99*, 103925.

Darity, William A. and Patrick L. Mason, "Evidence on Discrimination in Employment: Codes of Color, Codes of Gender," *Journal of Economic Perspectives*, Spring 1998, *12* (2), 63–90.

Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum, "The rise of the East and the Far East: German labor markets and trade integration," *Journal of the European Economic Association*, 2014, *12* (6), 1643–1675.

Deming, David and Lisa B. Kahn, "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals," *Journal of Labor Economics*, December 2017, *36* (S1), S337–S369.

Deming, David J and Kadeem Noray, "Earnings dynamics, changing job skills, and STEM careers," *The Quarterly Journal of Economics*, 2020, *135* (4), 1965–2005.

D'Haultfœuille, Xavier and Roland Rathelot, "Measuring segregation on small units: A partial identification analysis," *Quantitative Economics*, 2017, *8* (1), 39–73.

Dilger, Alexander and Lars Vischer, "No home bias in ghost games," *Discussion Paper of the Institute for Organisational Economics 7/2020*, 2020.
Diretta.it

Diretta.it, *Livesport Media Ltd, ST Balluta Business Centre, 10 Censu Tabone Street, St. Julians, STJ 1219, Malta.*

Dohmen, Thomas and Jan Sauermann, *"Referee bias," Journal of Economic Surveys, 2016, 30 (4), 679–695.*

Dohmen, Thomas J, *"Do professionals choke under pressure?," Journal of economic behavior & organization, 2008, 65 (3-4), 636–653.*

Doleac, Jennifer L and Benjamin Hansen, *"The unintended consequences of "ban the box": Statistical discrimination and employment outcomes when criminal histories are hidden," Journal of Labor Economics, 2020, 38 (2), 321–374.*

Dorn, David and Peter Levell, *"Trade and Inequality in Europe and the US," Institute for Fiscal Studies, 2021.*

Duncan, Otis Dudley and Beverly Duncan, *"A Methodological Analysis of Segregation Indexes," American Sociological Review, 1955, 20 (2), 210–217.*

Economist, The, *"Home comforts: Empty stadiums have shrunk home advantage, but not eliminated it," July 25th, 2020, page 73., 2020.*

_ , *"Non-white footballers played better when stadiums were empty during the pandemic," June 10th, 2021, www.economist.com/graphic-detail/2021/06/10/non-white-footballers-played-better-when-stadiums-were-empty-during-the-pandemic, 2021.*

Edmans, Alex, *"The link between job satisfaction and firm value, with implications for corporate social responsibility," Academy of Management Perspectives, 2012, 26 (4), 1–19.*

Efing, Matthias, Rudiger Fahlembrach, Christoph Herpfer, and Philipp Kruger, *"How Do Investors and Firms React to an Unexpected Currency Appreciation Shock?," Swiss Finance Institute Research Paper Series 15-65, Swiss Finance Institute December 2015.*

Endrich, Marek and Tobias Gesche, *"Home-Bias in Referee Decisions: Evidence from'Ghost Matches' During the COVID-19 Pandemic," Center for Law & Economics Working Paper Series, 2020.*
Fantacalcio.it

Fantacalcio.it, *QUADRONICA s.r.l. Via G. Porzio - CdN, Is. F4 80143, Napoli.*

Feri, Francesco, Alessandro Innocenti, and Paolo Pin, "Is there psychological pressure in competitive environments?," *Journal of Economic Psychology*, 2013, *39*, 249–256.

Ferraresi, Massimiliano and Gianluca Gucciardi, "Team performance and audience: experimental evidence from the football sector," Technical Report 2020.

Firpo, Sergio, Nicole M Fortin, and Thomas Lemieux, "Occupational tasks and changes in the wage structure," 2011.

Fischer, Kai and Justus Haucap, "Does crowd support drive the home advantage in professional soccer? Evidence from German ghost games during the COVID-19 pandemic," Technical Report, DICE Discussion Paper 2020.

Fitzpatrick, Thomas Bernard, "Soleil et peau," *Journal de Médecine Esthétiquer*, 1975.

Garicano, Luis, Ignacio Palacios-Huerta, and Canice Prendergast, "Favoritism under social pressure," *The Review of Economics and Statistics*, 2005, *87* (2), 208–216.

Garland, Jon and Michael Rowe, *Racism and anti-racism in football*, Springer, 2001.

Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card, "Assortative Matching or Exclusionary Hiring? The Impact of Firm Policies on Racial Wage Differences in Brazil," NBER Working Papers 25176, National Bureau of Economic Research, Inc October 2018.

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, April 2014, *104* (4), 1091–1119.

Gómez, Miguel A and Richard Pollard, "Reduced home advantage for basketball teams from capital cities in Europe," *European Journal of Sport Science*, 2011, *11* (2), 143–148.

Goos, Maarten, Alan Manning, and Anna Salomons, "Explaining job polarization: Routine-biased technological change and offshoring," *American economic review*, 2014, *104* (8), 2509–26.

_ **and** _ , "Lousy and lovely jobs: The rising polarization of work in Britain," *The review of economics and statistics*, 2007, *89* (1), 118–133.

Groshen, Erica L, "The structure of the female/male wage differential: Is it who you are, what you do, or where you work?," *Journal of Human Resources*, 1991, pp. 457–472.

Grossman, Gene M and Elhanan Helpman, "Product development and international trade," *Journal of political economy*, 1989, *97* (6), 1261–1283.

_ **and** _ , "Quality ladders and product cycles," *The Quarterly Journal of Economics*, 1991, *106* (2), 557–586.

_ **and** _ , "Quality ladders in the theory of growth," *The review of economic studies*, 1991, *58* (1), 43–61.

Grossmann, Sandra Hanslin, Sarah M. Lein, and Caroline Schmidt, "Exchange rate and foreign GDP elasticities of Swiss exports across sectors and destination countries," *Applied Economics*, 2016, *48* (57), 5546–5562.

Guardian, The, "Sharp rise in football racism as incidents go up by more than 50% in one year," *January 30th, 2020, www.theguardian.com/football/2020/jan/30/football-related-racist-incidents-sharp-rise-police-kick-it-ou*, 2020.

_ , "It's not just football. It's a barometer of diversity. Bring on the euros...," *June 6th, 2021, www.theguardian.com/football/2021/jun/06/its-not-just-football-its-a-barometer-of-diversity-bring-on-the-euros*, 2021.

Hanlon, W Walker, "Necessity is the mother of invention: Input supplies and directed technical change," *Econometrica*, 2015, *83* (1), 67–100.

Hellerstein, Judith, David Neumark, and Melissa McInerney, "Changes in Workplace Segregation in the United States between 1990 and 2000: Evidence from Matched Employer-Employee Data," in "The Analysis of Firms and Employees: Quantitative and Qualitative Approaches" NBER Chapters, National Bureau of Economic Research, Inc, December 2008, pp. 163–195.

Helleseter, Miguel Delgado, Peter Kuhn, and Kailing Shen, "The Age Twist in Employers' Gender Requests: Evidence from Four Job Boards," NBER Working Papers 22187, National Bureau of Economic Research, Inc April 2016.

_ **,** _ **, and** _ , "The Age Twist in Employers' Gender Requests: Evidence from Four Job Boards," *Journal of Human Resources*, October 2018, pp. 0416–7836R2.

Hersch, Joni and Jennifer Bennett Shinall, "Fifty years later: The legacy of the Civil Rights Act of 1964," *Journal of Policy Analysis and Management*, 2015, *34* (2), 424–456.

Hershbein, Brad and Lisa B. Kahn, "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings," *American Economic Review*, July 2018, *108* (7), 1737–1772.

Hijzen, Alexander, Holger Görg, and Robert C Hine, "International outsourcing and the skill structure of labour demand in the United Kingdom," *The Economic Journal*, 2005, *115* (506), 860–878.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang, "The wage effects of offshoring: Evidence from Danish matched worker-firm data," *American Economic Review*, 2014, *104* (6), 1597–1629.

Hylton, Kevin, "Black Lives Matter in Sport...?," *Equality, Diversity and Inclusion: An International Journal*, 2020.

Kaiser, Boris and Michael Siegenthaler, "The Skill-biased Effects of Exchange Rate Fluctuations," *The Economic Journal*, 2016, *126* (592), 756–780.

_ , _ , Andrin Spescha, and Martin Wörter, "The impact of real exchange rates on Swiss firms: Innovation, investment, productivity and business demography. Study on behalf of the State Secretariat for Economic Affairs SECO," Technical Report, KOF Studien 2018.

Kassimeris, Christos, *European football in black and white: tackling racism in football*, Lexington Books, 2007.

Kaufmann, Daniel and Tobias Renkin, "Manufacturing prices and employment after the Swiss franc shock," Strukturberichterstattung, State Secretariat for Economic Affairs SECO October 2017.

_ and _ , "Export Prices, Markups, and Currency Choice after a Large Appreciation," IRENE Working Papers 19-07, IRENE Institute of Economic Research October 2019.

Kettemann, Andreas, Andreas I Mueller, and Josef Zweimüller, "Vacancy durations and entry wages: evidence from linked vacancy-employer-employee data," Technical Report, National Bureau of Economic Research 2018.

Kinlaw, Dennis C, *Coaching for commitment: Interpersonal strategies for obtaining superior performance from individuals and teams*, Jossey-Bass/Pfeiffer, 1999.

Kline, Patrick, Evan Rose, and Christopher Walters, "Systematic Discrimination Among Large U.S. Employers," NBER Working Paper 29053, National Bureau of Economic Research 2021.

Koning, Ruud H, "Home advantage in professional tennis," *Journal of Sports Sciences*, 2011, *29* (1), 19–27.

Krusell, Per, Lee E Ohanian, José-Víctor Ríos-Rull, and Giovanni L Violante, "Capital-skill complementarity and inequality: A macroeconomic analysis," *Econometrica*, 2000, *68* (5), 1029–1053.

Kuhn, Peter and Kailing Shen, "Gender Discrimination in Job Ads: Evidence from China," *The Quarterly Journal of Economics*, 2013, *128* (1), 287–336.

Kuhn, Peter J and Kailing Shen, "What Happens When Employers Can No Longer Discriminate in Job Ads?," Technical Report, National Bureau of Economic Research 2021.

Kuhn, Peter, Kailing Shen, and Shuo Zhang, "Gender-targeted job ads in the recruitment process: Facts from a Chinese job board," *Journal of Development Economics*, 2020, *147*, 102531.

Liardi, Vincent L and Albert V Carron, "An analysis of National Hockey League face-offs: Implications for the home advantage," *International Journal of Sport and Exercise Psychology*, 2011, *9* (2), 102–109.

Lujanski-Lammer, Elke, "Stellenausschreibungen – ein Spiegelbild der Gesellschaft [Job Advertisements – A Mirror Image of Society]," Diplomarbeit, University of Graz, Institute for Pedagogy 2006.

Mihaylov, Emil and Kea Gartje Tijdens, "Measuring the Routine and Non-Routine Task Content of 427 Four-Digit ISCO-08 Occupations," 2019.

Mirkov, Nikola, Igor Pozdeev, and Paul Söderlind, "Toward removal of Swiss franc cap: Market expectations and verbal interventions," 2016.

Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance, "Downskilling: changes in employer skill requirements over the business cycle," *Labour Economics*, August 2016, *41*, 333–347.

Mueller, Andreas I., Damian Osterwalder, and Josef Zweimueller, "The Austrian AMS Vacancy Database," Mimeo, University of Zurich August 2019.

Mumford, Karen and Peter N. Smith, "What Determines the Part-Time and Gender Earnings Gaps in Britain: Evidence from the Workplace," *Oxford Economic Papers*, 2009, *61*, i56–i75.

Oswald, Andrew J, Eugenio Proto, and Daniel Sgroi, "Happiness and productivity," *Journal of labor economics*, 2015, *33* (4), 789–822.

Pagán-Castaño, E, A Maseda-Moreno, and C Santos-Rojo, "Wellbeing in work environments," *Journal of Business Research*, 2020, *115*, 469–474.

Palacios-Huerta, Ignacio, *Beautiful game theory: How soccer can help economics*, Princeton University Press, 2014.

Petersen, Trond and Laurie A Morgan, "Separate and unequal: Occupation-establishment sex segregation and the gender wage gap," *American Journal of Sociology*, 1995, *101* (2), 329–365.

Pettersson-Lidbom, Per and Mikael Priks, "Behavior under social pressure: Empty Italian stadiums and referee bias," *Economics Letters*, 2010, *108* (2), 212–214.

Pollard, Richard, "Home advantage in soccer: variations in its magnitude and a literature review of the inter-related factors associated with its existence," *Journal of Sport Behavior*, 2006, *29* (2), 169.

Reade, James, Dominik Schreyer, and Carl Singleton, "Echoes: what happens when football is played behind closed doors?," *Available at SSRN 3630130*, 2020.

_ , _ , and _ , "Eliminating supportive crowds reduces referee bias," *Available at SSRN 3743972*, 2020.

Scoppa, Vincenzo, “Social pressure in the stadiums: Do agents change behavior without crowd support?,” *Journal of Economic Psychology*, 2021, *82*, 102344.

Shields, Michael A and Stephen Wheatley Price, “Racial harassment, job satisfaction and intentions to quit: evidence from the British nursing profession,” *Economica*, 2002, *69* (274), 295–326.
Soccerwiki.org

Soccerwiki.org, *Soccer Manager Ltd, 4th Floor 33 Park Place, Leeds, West Yorkshire, England, LS1 2RY.*

Sors, Fabrizio, Michele Grassi, Tiziano Agostini, and Mauro Murgia, *“The sound of silence in association football: Home advantage and referee bias decrease in matches played without spectators,” European Journal of Sport Science, 2020, pp. 1–21.*

Stoermer, Sebastian, Azusa Hitotsuyanagi-Hansel, and Fabian Jintae Froese, *“Racial harassment and job satisfaction in South Africa: The moderating effects of career orientations and managerial rank,” The International Journal of Human Resource Management, 2019, 30 (3), 385–404.*

SWISSMEM, *“Solution provider for the challenges of our time. The Swiss MEM industries,” 2019.*

Switzerland Global Enterprise, *“Advanced Manufacturing in Switzerland. Switzerland Global Enterprise,” 2020.*

Tanaka, Mari, *“Exporting Sweatshops? Evidence from Myanmar,” The Review of Economics and Statistics, 2020, 102 (3), 442–456.*

Thoenig, Mathias and Thierry Verdier, *“A theory of defensive skill-biased innovation and globalization,” American Economic Review, 2003, 93 (3), 709–728.*
Transfermarkt.com

Transfermarkt.com, *Transfermarkt GmbH & Co. KG, Wandsbeker Zollstrasse 5a, 22041 Hamburg.*

Verhoogen, Eric A, “Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector,” *The Quarterly Journal of Economics*, 2008, *123* (2), 489–530.

Wozniak, Abigail, “Discrimination and the effects of drug testing on black employment,” *Review of Economics and Statistics*, 2015, *97* (3), 548–566.

Ziegler, Alexander, “Skill Demand and Wages. Evidence from Linked Vacancy Data,” IZA Discussion Papers 14511, Institute of Labor Economics (IZA) 2021.

Zweimüller, Josef, Rudolf Winter-Ebmer, Rafael Lalive, Andreas Kuhn, Jean-Philipe Wuellrich, Oliver Ruf, and Simon Büchi, “Austrian Social Security Database,” 2009.