Essays on Technology, Trade and Labour Markets



Katherine Stapleton

Lincoln College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy in Economics

August 2020

Essays on Technology, Trade and Labour Markets

Katherine Stapleton

Lincoln College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Economics

August 2020

This thesis consists of three standalone chapters. Chapter 1 investigates how automation in high-income countries affects trade and multinational activity involving less developed ones. This chapter studies Spain's manufacturing sector from 1990 to 2016 and exploits advances in the capabilities of robots over time, finding that the use of robots in Spanish firms actually had a positive impact on their imports from, and number of affiliates in, less developed countries. We provide a modelling framework and empirical analysis on why this has been the case. Chapter 2 explores how trade policy uncertainty and the threat of a reversal of deep integration affects labour demand. This chapter studies how the the threat of trade barriers caused by the Brexit referendum affected the posting of online job adverts between 2015 and 2019, finding that local labour markets that were more exposed to the threat of services trade barriers experienced a relative decline in online job adverts after the referendum, particularly in months when there was heightened uncertainty about trade policy, but the threat of tariffs on goods had little impact. Chapter 3 is on a similar theme to Chapter 1 and explores how the use of Artificial Intelligence (AI) in high income countries affects services offshoring to less developed countries. I construct new measures of the demand for Machine Learning (ML) skills in the UK based upon the text of online job adverts. Using these measures and instrumenting for ML demand using advances in task-specific capabilities of AI over time, I show that the deployment of ML in the UK has had a positive impact on services offshoring, particularly to low and middle income countries.

Supervisors This thesis was supervised by Chris Woodruff and Beata Javorcik.

Total words This thesis contains approximately 71,719 words.

Statement of authorship Chapter 1 is joint work with Michael Webb. Chapter 2 is joint work with Beata Javorcik, Ben Kett and Layla O'Kane.

Acknowledgement of funding I acknowledge funding from the Economic and Social Research Council (ESRC) and the Sloane Robinson Foundation.

Acknowledgements

I was warned prior to starting that doing a PhD in Economics can be a lonely and arduous process. The fact that I have been pleasantly surprised by how much I have enjoyed the past four years is a credit to the wonderful supervisors, peers and research community that I have had the privilege of working with in Oxford, all of whom I thank deeply and unreservedly.

To my supervisors, Beata Javorcik and Chris Woodruff: I am deeply grateful for your support, encouragement and the wisdom and knowledge you have imparted on me. I have been consistently blown away by your dedication, reliability and generosity with your time; I only hope I have managed to absorb enough. To my initial supervisors at the beginning of the DPhil, Doug Gollin and Kalina Manova: thank you for your help and encouragement getting me started on the DPhil. To Nick Bloom, my supervisor when visiting at Stanford: thank you for being an excellent host, for your feedback and all that you and the Stanford research community taught me.

To the inspiring academics I have had the privilege of learning from in Oxford: the remaining faculty in the Center for the Study of African Economies (CSAE), Stefan Dercon, Simon Quinn, Pramila Krishnan, Chris Adam and Gabriel Ulyssea: thank you for teaching me much of what I know in the field of development economics, for your kindness and for your feedback. More importantly, thank you for making CSAE an academic 'home' for me in Oxford. Having this remarkably friendly research community has made doing a PhD an infinitely happier experience. In the trade group, Peter Neary, Ferdinand Rauch, Tony Venables and Kevin O'Rourke: thank you for imparting on me so much knowledge about international trade and inspiring and developing my ideas.

To the wonderful individuals with whom I started the DPhil journey (or who started just after me) and who subsequently became the best of friends and coauthors, Ashley Pople, Ben Kett and Alex Copestake: thank you for helping me at all the hard times and for your fun, your ideas, your advice and your friendship in all the good times - if I had not started with such great peers, I am not sure I would have even made it through the first year. To the others that made my DPhil days so much brighter, Vanessa, Verena, Hannah, Maria, Simon and Lukas: thank you for the dinners, the discussions, the learning and the support. And to my other coauthors, Michael Webb and Friedrich Bergmann: thank you for being great collaborators and a pleasure to work with.

To the others in Oxford who made these past few years so formative and memorable, Almut, Andre, James, Kathy, Marta, Audrey, Ruby, Shazea, Wanyu, Diva, Ambika: thank you for your support and friendship at all different stages of my Oxford life. To other friends far and wide, Liz: thank you for keeping me on track, and for taking me on hiking trips when Hogwarts was getting too much, Arendse: thank you for spending so many days working with me in cafes in Beijing when I was writing this thesis, Leon and Kai: thank you for motivating me on the economics quest many years back and for the spontaneous trips, and to my dear friends from home, Madhu, Beth, Becky, Romy, Lauren, Steph, Erin, Harmeet and Chloe: thank you for your inspiration and your unwavering friendship.

To my family, my parents, Maire and Chris: thank you for always encouraging me in my education, for inspiring me to be curious and follow my interests and for supporting me in all of my pursuits. To James, Emily and Eleanor: thank you for reminding me to work less and spend more time climbing mountains or riding bikes. To my granddad: thank you for being a role model on the value of determination and perseverance. And last, but not least, to Mike: thank you for cheering me along at each step of the way during the writing of this thesis.

I am deeply grateful to the funding I received from the ESRC, Lincoln College, the Economics Department and the grants from CSAE for the data for Chapter 1 and exploratory fieldwork in relation to Chapter 3. Without these, I would not have had the privilege of being able to undertake this research or conduct research visits, attend conferences or learn about the real world context behind my research. I have also been extremely fortunate to work with Burning Glass Technologies on two of my PhD projects. To Bledi Taska and Layla O'Kane at BGT: I thank you wholeheartedly for being so helpful and for making this aspect of my PhD so hassle free and enjoyable.

Katherine Stapleton
Oxford, August 2020

Contents

Introduction **1**

1 Automation, Trade and Multinational Activity: Micro Evidence from Spain **9**
1.1 Introduction . . . 10
1.2 Automation and offshoring with heterogeneous firms . . . 15
1.3 Data . . . 22
1.4 Which firms automate? . . . 35
1.5 Empirical strategy . . . 42
1.6 Results . . . 50
1.7 Robustness . . . 60
1.8 Conclusion . . . 67

Appendices of Chapter 1 **71**
1.A Additional figures . . . 72
1.B Additional tables . . . 75
1.C Additional definitions . . . 81

2 Unravelling Deep Integration: Local Labour Market Effects of the Brexit Vote **82**
2.1 Introduction . . . 83
2.2 Background on the Brexit Referendum . . . 90
2.3 Data . . . 95
2.4 Empirical strategy . . . 115
2.5 Results . . . 116
2.6 Robustness checks . . . 130
2.7 Conclusion . . . 136

Appendices of Chapter 2 **141**
2.A Comparison of BGT data with other sources . . . 142
2.B Additional maps and figures . . . 146
2.C Additional tables . . . 149

3 Artificial Intelligence and Services Offshoring **159**
3.1 Introduction . . . 160
3.2 Background . . . 166
3.3 Data . . . 169
3.4 Measuring demand for ML skills . . . 179
3.5 Empirical strategy . . . 183
3.6 Results . . . 188
3.7 Robustness . . . 200
3.8 Conclusions . . . 207

Appendices of Chapter 3 **211**
3.A Additional Definitions . . . 212

Conclusion **213**

References **216**

List of Figures

1.1 Key domestic outcomes over time 24
1.2 Share of firms using each technology, by industry 25
1.3 Key international outcomes over time 27
1.4 Mean less developed import intensity, share with affiliates & no. of affiliates, by industry 29
1.5 Sequencing of automation and importing intensively from less developed countries 39
1.6 Labour productivity distribution by automation and importing from less developed countries 41
1.7 Top 10 industry-region pairs with highest and lowest median exposure 47
1.8 Technology use by region 72
1.9 Sequencing of automation and importing intensively from less developed countries 72
1.10 Labour productivity distribution by automation and less developed affiliate status 73
1.11 Automation, importing intensively or affiliates in less developed countries, by industry 74

2.1 Monthly online job postings, UK and US 84
2.2 Monthly online job postings in the UK 97
2.3 Country-sector pairs with highest STRI 'gap' 101
2.4 Exposure by Travel to Work Area: professional services trade barriers 103
2.5 Most exposed Travel to Work Areas: professional services trade barriers 103
2.6 Exposure by Travel to Work Area: MFN tariffs 106
2.7 Most exposed Travel to Work Areas: MFN tariffs 106
2.8 Skilled and unskilled job postings over time 108
2.9 Brexit trade policy uncertainty measures 111
2.10 Impact of professional services exposure over time 118
2.11 Impact of financial services exposure over time 128
2.12 Average monthly vacancies: BGT and UK Vacancy Survey 143

2.13 BGT average monthly vacancies and ONS average monthly jobs . . . 145
2.14 UK total trade in services in £millions 146
2.15 Financial service exposure map 147
2.16 Most exposed TTWAs: Financial Services 147
2.17 Exposure to the exchange rate depreciation and immigration 148

3.1 UK service imports from top 4 LMI country partners 175
3.2 UK service imports from LMI countries by product 176
3.3 ML-related skills 180
3.4 Top 5 industries by demand for ML skills 182
3.5 Geography of ML hiring 183
3.6 Occupations with highest exposure scores 185
3.7 Correlation between establishment baseline ML exposure and change in demand for ML skills 190

List of Tables

1.1 Summary statistics . . . 32
1.2 Prevalence & market share of firms that automate . . . 35
1.3 Technology adoption premia . . . 36
1.4 Future technology adoption premia . . . 38
1.5 Baseline IV results . . . 52
1.6 Extensive margin and affiliate activities . . . 54
1.7 Sequencing of importing intensively and robot use . . . 56
1.8 Changes to employment and operations in Spain . . . 58
1.9 Intensive margin of automation . . . 59
1.10 Controlling for changes to world export supply . . . 61
1.11 Controlling for import tariff changes . . . 62
1.12 Excluding the financial crisis . . . 63
1.13 IFR instrument for robot adoption . . . 64
1.14 Excluding the automotive sector . . . 66
1.15 Alternative specifications for imports regressions . . . 67
1.16 Correlation between technology use . . . 75
1.17 ISCO 68 codes with highest & lowest patent scores . . . 75
1.18 Results without IV . . . 76
1.19 Baseline results excluding firms with <5 year obs . . . 77
1.20 Baseline IV results: weighted . . . 78
1.21 Changing thresholds for import intensity . . . 79
1.22 Olley-Pakes TFP . . . 80
1.23 Countries that join OECD & EU during sample period . . . 81

2.1 Brexit Timeline . . . 92
2.2 Summary statistics . . . 95
2.3 Posting occupations and skill levels . . . 107
2.4 Uncertainty measure included terms . . . 110
2.5 Baseline post-vote results . . . 117
2.6 Uncertainty measure results . . . 121
2.7 US Placebo Test . . . 122

2.8 Impact by skill group . . . 124
2.9 Professional service results breakdown by sector . . . 125
2.10 Financial services post-vote baseline . . . 127
2.11 Financial services impact by skill group . . . 129
2.12 Financial services impact by occupation . . . 131
2.13 Mapping between UK SIC, OECD STRI & ONS Service export product types . . . 149
2.14 US placebo for Financial Services . . . 149
2.15 Uncertainty measures with quarterly trend controls . . . 150
2.16 Professional service exposure impact by occupation . . . 151
2.17 Baseline Excluding London . . . 152
2.18 Including shares . . . 153
2.19 Alternative tariff measures . . . 154
2.20 Goods non-tariff barriers . . . 155
2.21 Specific Tariffs and Quotas . . . 156
2.22 Import protection and imported inputs . . . 157
2.23 Pre-vote placebo for 2015 . . . 158

3.1 Posting occupations and skill levels . . . 173
3.2 Summary statistics . . . 178
3.3 Firm level results: Impact of ML exposure on growth in ML skills demand . . . 189
3.4 Firm level IV results: impact of ML demand on non-ML & total postings . . . 191
3.5 Firm level IV results: impact of ML demand on non-ML postings by skill type . . . 192
3.6 Firm level IV results: impact of ML demand on most & least exposed occupations . . . 194
3.7 TTWA IV results: Impact of ML demand on services offshoring . . . 195
3.8 TTWA IV results: impact of ML demand on offshoring by LMI country & service type . . . 197
3.9 Industry IV results: Impact of ML demand on offshoring . . . 198
3.10 Industry IV results: impact of ML demand on offshoring by country & service type . . . 199
3.11 Firm level IV results: replacing ΔML share with ΔIHS(ML postings) 201
3.12 Firm level IV results: increasing threshold to 20 . . . 202
3.13 Firm level IV results: 2012-2015 . . . 203
3.14 TTWA level IV results: replacing ML share with IHS(ML postings) . 204
3.15 TTWA level IV results controlling for computer service expenditure . 206

Introduction

This thesis presents research on two key topics relating to recent changes affecting international trade and global supply chains. The first topic is how technological advances, primarily those in robotics and AI, affect trade and multinational activity, and particularly, trade with developing countries. Chapters 1 and 3 both present research on this topic. These two papers investigate how automation in the manufacturing sector, and AI in the services sector, respectively, of high income countries, affect offshoring to less developed countries. The second topic is how the threat of trade barriers and a reversal of global trade integration affects labour markets. Chapter 2 aims to provide evidence on this question in the context of the Brexit referendum, by studying how the threat of trade barriers on UK exports to the EU affected the posting of online job adverts.

My initial motivation to conduct research on the first topic stemmed from professional experience attempting to design export promotion and FDI attraction policies aimed at increasing growth in the manufacturing sector in Rwanda. The historical importance of manufacturing and industrialisation as drivers of economic growth and poverty alleviation have been widely acknowledged. Over the last four decades, offshoring and the global fragmentation of value chains have played a major role in enabling this path of manufacturing-led development. In light of the development success of countries that followed this path, particularly in East Asia, substantial policy efforts and vast sums of money have since been dedicated, in other developing countries like Rwanda, to trying to attract manufacturing investment and become integrated into Global Value Chains (GVCs). For example, between 1990 and 2009, the global number of national export promotion agencies tripled and between 1995 and 2018, around 5000 Special Economic Zones have been built around the world (Lederman et al., 2009; UNCTAD, 2019).

It has often been assumed that manufacturing production will inevitably next relocate to a new wave of developing countries in South-East Asia and Africa. For example, in 2011 there was talk of the '85 million jobs' that will be imminently moving out of China (Lin, 2011). Against this backdrop, however, there have been a number of changes to global trade and the nature of value chains over the past decade. After rising substantially for two decades, the growth of a variety of measures of GVC trade relative to GDP fell sharply in 2008 and have since levelled out (Li et al., 2019). This has caused some debate over whether the process of deeper GVC integration has stalled or started to reverse. Value chains are becoming more regionally concentrated and the share of global trade between high- and low-income countries has been declining (Stapleton, 2019).

At the same time, industrial policies implemented in parts of sub-Saharan Africa (SSA) appear to have had mixed results. Manufacturing value added as a share of GDP for SSA was lower in 2018 than 2000, at 10.6% relative to the global average of 16.8% (World Bank, 2020). While the topic is much debated, there have been concerns that SSA is de-industrialising. These facts all spurred my curiosity to understanding in more depth the extent to which global production activities were being relocated to new destinations, the factors determining firm decisions surrounding when to offshore and where to locate different stages of production and the recent changes to GVCs since the financial crisis.

There has been much discussion, particularly in the media and in policy institutions, about the impact of automation and emerging technologies on manufacturing production, global supply chains and the scope for manufacturing-led development. It is regularly assumed that the automation of manufacturing processes will reduce the incentive of firms to relocate production to areas with lower labour costs. An often discussed hypothesis is that automation may lead firms to 'reshore' production and return it to where it was previously conducted, typically assumed to be a high-income country. There have been a number of high-profile anecdotal examples used as evidence of such a trend. For example, in 2015 Adidas received much attention for establishing its 'Speedfactories' in Germany and the US, to produce athletic footwear using robotic cutting and 3-D printing.

Another, less frequently discussed, hypothesis has been that manufacturing

production will instead remain concentrated in existing production locations, such as China, where it will be automated instead of relocated. A regularly cited example is that of Foxconn, a contract manufacturer best known for producing Apple and Samsung products in China, which replaced 60,000 Chinese factory workers with industrial robots in 2016 (South China Morning Post, 2016). Proponents of this view point out that over the past decade China has experienced one of the world's fastest growth rates in installations of industrial robots and offers substantial government support for the robotics industry and private sector uptake of automation (IFR, 2019; Cheng et al., 2019).

Despite the media and policy attention, however, I was surprised to discover that beyond these anecdotal case studies, there is very little systematic empirical evidence on this topic in the economics literature. Automation in manufacturing and particularly, industrial robots, have been used commercially for a number of decades now, yet there has been very little research on how they have affected trade, global supply chains or offshoring. In fact, more generally, we know remarkably little about the detailed micro-level impacts of the adoption of different technologies that automate production. Technology has often been treated as a black box in economics and detailed data on firm use of specific machines or technologies is scarce; this has been pointed out by Seamans and Raj (2018), who note that the scarcity of firm-level data has hindered our ability to understand the impacts or automation, or as Seamans (2017) has put it in his aptly titled article, 'We Won't Even Know If A Robot Takes Your Job'. This dearth of detailed understanding on the topic of automation spurred me to search for firm -level data with details on both detailed technology use and trade or global supply chain activities, and subsequently led to the writing of Chapter 1 of this thesis.

Chapter 1, written with Michael Webb, makes use of a relatively unique dataset of Spanish manufacturing firms that provides both rich information on firm technology use, covering three specific technologies related to automation: robots, computer numerically controlled (CNC) machines and flexible manufacturing systems (FMS) and also details on trade and multinational activity. We begin by documenting the characteristics of firms that have adopted these three technologies between 1990 and 2016 and how technology adoption relates to international trade

and multinational activity. We explore the sequencing of firm decisions regarding importing, opening affiliates and automation and the productivity distribution of firms that automate, offshore or do both. Automation is unlikely to be exogenous to firm decisions regarding trade and multinational activity. To evaluate the causal impact of automation we therefore exploit supply-side advances in specific capabilities of robots over time, as documented in the text describing robotics patents, that have made it technically feasible to automate some tasks and not others.

In this paper we find that fears of the consequences of automation for reshoring seem to have been overblown and have over-simplified what is, in fact, a very complex relationship. In our sample of firms in Spain, more firms started to automate before they started to import intensively from, or open affiliates in, less developed countries. For these firms, automation generated new demand for imports from less developed countries. For firms that offshored before they starting to use robots, robot adoption had no effect on the value of imports, but decreased the share of imports sourced from developing countries. The former effect dominates and for the full sample robot use has a positive effect on imports from, and the number of affiliates in, less developed countries.

More generally, we find substantial heterogeneity in terms of technology adoption. The existing literature has typically studied and modelled automation at the industry level, missing a lot of these important dynamics and details. We try to explain our findings by extending the existing canonical task-based framework for automation to include the two additional ingredients of firm heterogeneity and the option to offshore production. This work scratches the surface of a broad topic in terms of what we are able to observe and conclude with this data; in the conclusion of this thesis I discuss the unanswered questions that remain and avenues for further work.

Concerns surrounding the feasibility of manufacturing-led-development as a growth path tend to lead to a discussion about an alternative services led-development model. There are several countries that have experienced productivity-enhancing structural transformation through the reallocation of economic activity from agriculture directly into services. In India particularly, the positive contribution of structural change to economic growth after the 1990s was

largely down to the expansion of high-productivity export-oriented service activities such as IT, business process outsourcing (BPO), and other business services (McMillan et al., 2017). The features of the manufacturing sector that were conducive to development, such as the capacity to employ large numbers of low skilled workers at a productivity premium and the potential for technology diffusion and spillover effects, may well be increasingly exhibited by some of these export-oriented service industries. This has led to increased focus on the potential for these tradable service sectors to be a future growth path for developing countries (see, for example, Baldwin and Forslid (2020)).

Once again, technological advances could be expected to play a defining role in the evolution of services supply chains, yet empirical evidence to date is limited. There is a growing literature suggesting that improvements in communications technology enabled the fragmentation of GVCs and hence paved the way for the offshoring of services activities to countries like India and the Philippines (e.g. Fort (2017) , Steinwender (2018)). There have been suggestions that advances in AI and digital platforms, such as UpWork and Amazon MTurk, may enable a greater degree of offshoring (see, for example Baldwin (2019)). There is already suggestive evidence that machine translation has a positive enabling effect on trade (Brynjolfsson et al., 2019). In contrast, there has also been speculation that ML could reduce the offshoring of many services activities, particularly repetitive back-office activities, such as data entry, invoice processing and payroll, which can increasingly be automated due to advances in some subfields of ML. Curiosity about the impact of AI on services offshoring spurred my motivation to write Chapter 3.

The impacts of AI on labour markets and trade are likely to be more complex than those of robots. While robots are a simple automation technology used for a narrow set of tasks, there is growing evidence that AI is a General Purpose Technology (GPT), one with potential for pervasive use in a wide range of sectors and the ability to catalyse new waves of innovation (Cockburn et al., 2018). Empirical evidence on the impacts of AI is even more scarce than for robots, not least because the recent wave of advances in AI happened only in the past decade. I was therefore motivated to find a way to identify new ways of measurement for the deployment of AI and took inspiration from Tambe and Hitt (2012) and Rock

(2019) to backwards induce the deployment of AI and specifically, ML, from the demand for ML skills. I do this in Chapter 3 using a dataset of the majority of all job adverts posted online in the UK between 2012 and 2017, scraped and collated from a large number of websites by Burning Glass Technologies (BGT), a private market research firm. I use this data to develop proxy measures of the deployment of ML based upon the reported demand for ML skills in the job adverts. I then combine these measures with data from the UK's Inquiry on International Services (ITIS), which provides detailed information on UK firm imports of services by product type and country, allowing me to construct detailed measures of services offshoring.

I exploit plausibly exogenous changes in the tasks that have become feasible to automate over time and show that firms, regions and industries with more potential to deploy ML to automate their tasks have experienced a greater rise in demand for ML skills between 2012 and 2017. The deployment of ML itself has had a positive impact on offshoring, particularly to low and middle income countries, rather than high income countries. Within this group, the main beneficiary has been India, the UK's largest services offshoring destination. This impact has been primarily for the offshoring of business services, which includes financial services, insurance, legal and accounting, amongst others, and ICT-related services, comprised of information services, computer services and telecommunication services. Therefore despite AI being a very different technology to robots, these results are quite similar to what was found in Chapter 1, that the productivity or enabling effect of AI may outweigh the displacement effects.

Chapter 2 is on the separate theme of the second question. The past four years when I have been writing my DPhil have been characterised by a backlash against the heightened global trade integration of the preceding decades. While there is a substantial body of evidence in the trade literature about the impacts of trade integration, the removal of tariffs and trade liberalisation, there is relatively scarce evidence on the consequences of a move in the opposite direction and the reversal of existing trade agreements. The Brexit vote has posed the threat of a reversal of trade integration on an unprecedented scale: the potential unravelling of decades worth of deep integration with the world's most advanced and integrated trad-

ing bloc. In addition, the UK is a largely services dominated economy and has complex and advanced services supply chains that have been governed by deep provisions in diverse areas, such as intellectual property rights, investment and product standards. These are aspects of trade policy that are generally under-researched and we have relatively little evidence on how they affect trade and labour markets. These factors formed my motivation for wanting to understand the consequences of this scenario of a reversal of trade integration for workers and individuals, through studying the impact on labour demand.

In Chapter 2, written with Beata Javorcik, Ben Kett and Layla O'Kane, we use the same dataset of job adverts collected by BGT as in Chapter 3, to explore how the Brexit vote affected UK labour demand, proxied by the posting of job adverts, throughout the Brexit negotiation period. We exploit pre-referendum variation in EEA country-sector MFN service trade restrictions relative to restrictions applicable to non- EEA countries to develop measures of UK local labour market exposure to the threat of no-deal scenario export barriers for professional services exports. We construct an analogous measure for goods exports based on EEA MFN tariffs. We find that the threat of trade barriers on professional services exports to the EU had a strong negative effect on the posting of online job adverts after the Brexit referendum, particularly for higher-skilled jobs. In contrast to services barriers, the threat of future tariffs on goods does not seem to have impacted the posting of job adverts. We conclude that even the threat of unravelling deep integration with a country's major trading partners can have substantial effects on labour markets.

Contributions

The research in this thesis makes a number of contributions to the literatures on international trade, labour and development economics. Chapter 1 provides three key contributions to the automation literature and the literature on how emerging technologies impact trade. The first and main contribution is to evaluate the causal impact of the use of robots on trade and multinational activity at the firm level, contributing to a literature that has typically evaluated this relationship at the industry level. The second key contribution is to document the firm heterogeneity and detailed characteristics of manufacturing firms that automate using

three distinct technologies and to provide descriptive details on the relationship between firm technology adoption and international orientation. The final key contribution is to offer the building blocks of a modelling framework that is able to explain some of the key aspects of this relationship.

Chapter 2 provides the primary contribution of investigating the impact of an event of major trade policy uncertainty, and particularly the threat of trade barriers, on labour demand as measured using the posting of job adverts, contributing to a literature that has so far mainly focused on how trade policy uncertainty affects trade and investment alone. This chapter additionally provides new evidence not only on the threat of barriers to trade in goods, but also of services trade barriers specifically, which have so far received relatively little attention. Finally, this chapter provides the contribution of documenting the substantial decline in the posting of job adverts in the UK after 2017 and offers evidence to explain that decline.

In Chapter 3 I provide a number of contributions to the emerging literatures on the economic impacts of AI and on services trade. The first contribution is one of measurement: I develop a new way to measure AI deployment by backwards induction from the skills being demanded in job adverts. This contributes to a literature where, to date, we have very little information on whether, and to what extent, firms and industries are actually using ML commercially. The second key contribution is to evaluate the causal impact of ML deployment on services offshoring, particulary to low and middle income countries. Finally this chapter also contributes to the literature on the impacts of AI on domestic labour markets, which to date has been scarce, by empirically evaluating the causal impact of ML deployment on labour demand in the UK.

CHAPTER 1

Automation, Trade and Multinational Activity: Micro Evidence from Spain

with **Michael Webb**

Abstract We use a rich dataset of Spanish manufacturing firms from 1990 to 2016 to shed new light on how automation in a high-income country affects trade and multinational activity involving less developed countries. By exploiting supply side improvements in the capabilities of robots over time, as described in patents, we show that, contrary to the speculation that automation will cause reshoring, the use of robots in Spanish firms actually had a positive impact on their imports from, and number of affiliates in, less developed countries. Robot use causes firms to expand production, increase productivity and increases the probability that they start importing from, or opening affiliates in, less developed countries. The sequencing of automation and offshoring has important consequences for the impact of automation, however. For firms that offshored before they starting to use robots, robot adoption had no effect on the value of imports, but decreased the share of imports sourced from developing countries. By contrast, for firms that had not already offshored, robot adoption made them more likely to start doing so. We show that these findings can be explained in a framework that incorporates firm heterogeneity, the choice between automation, offshoring and performing tasks at home and where automation and offshoring both involve upfront fixed costs, such that their sequencing matters.

1.1 Introduction

Recent technical advances in robotics and machine intelligence have been accompanied by a new wave of concern about the consequences of automation. A growing body of literature has been dedicated to examining the impacts of emerging labour-replacing technologies, particularly industrial robots. This research has typically focused on the impacts on domestic labour markets in high-income countries, in isolation from trade patterns and global supply chains (for example, Acemoglu and Restrepo (2019a); Graetz and Michaels (2018)). Yet, in a globally integrated economy, it might be expected that automation in one country will not only affect its own labour markets, but also those of its trading partners. Over the past few decades, high-income countries experienced a substantial wave of offshoring of labour intensive manufacturing activities to less developed countries. In a world where production is already conducted offshore, automation in a high-income country could have an important effect on trade patterns and labour markets in offshore production destinations.

This is a concern that has been emanating from policy institutions and development agencies: that automation technologies might substitute for low-skilled labour in less developed countries, thus reducing the future scope for manufacturing-led development in parts of the world that have yet to industrialise.[1] Manufacturing-led development has been an unrivalled force for poverty reduction. Between 1990 and 2015, 1.2 billion people were lifted out of poverty and of this decline, 93% came from a select group of countries that became heavily integrated in global supply chains over this period – China, India, Korea, Mexico, Poland, Indonesia and Thailand.[2] This question therefore has important implications, yet there is relatively little empirical evidence to date on whether this fear is based on reality, or even how advances in labour-replacing technologies over the past three decades have affected trade and global supply chains involving developing countries.

A key part of understanding the relationship between technological advances that make it easier, or cheaper, to automate production and offshoring is to con-

[1] See, for example, Hallward-Driemeier and Nayyar (2017)

[2] Based on the World Bank's US$1.9-a-day poverty line. Poverty data from World Bank Povcalnet http://iresearch.worldbank.org/PovcalNet/data.aspx

sider firm level decisions on whether to offshore production stages and whether to invest in labour-replacing technology. In a world of advanced technology and global supply chains, manufacturers increasingly have access to a range of strategies to reduce their production costs, become more competitive or increase quality. The choice between these strategies depends on their relative costs and benefits. Existing research has tended to study automation at the industry level, missing out on these key firm decisions. To date we, in fact, have surprisingly little empirical evidence about why, and when, firms automate and how it affects firm outcomes, domestic or international.

In this paper we show that firm heterogeneity and dynamics are important for understanding the consequences of automation. We take advantage of a unique dataset of Spanish manufacturing firms that provides both rich information on firm technology use and details on trade and multinational activity between 1990 and 2016. It is well known in the international trade literature that only a small subset of firms export, import or offshore (Bernard et al., 2003, 2009). We show that a similar pattern holds for technology adoption: only a small subset of firms adopt three distinct automation technologies: robots, computer numerically controlled (CNC) machines and flexible manufacturing systems (FMS), but these firms account for the majority of employment, production and trade. For example, in 2014, 33% of the firms in our sample used robots, but these firms accounted for 64% of manufacturing employment, 66% of value added and 72% of exports. Firms that automate also have distinct characteristics: they are on average around twice as large, have around 60% higher TFP, export and import nearly four times more, are around 10% more likely to be a multinational and have an approximately 5% lower labour share than other firms in the same industry, region and year.

We find that prior to having adopted these technologies, firms that will automate in the future already import more and have more affiliates in all regions, including less developed countries, but import a lower share from less developed countries. This gap is then wider, in any time period, between firms that have automated and those that haven't. Contrary to the typical assumption that automation might cause 'reshoring', we find that firms in our sample were, in fact, more likely to have started automating, using either robots, FMS or CNC machines, before

starting to import intensively from, or open affiliates in, less developed countries. When considering robots alone, firms were about as likely to have started using robots before importing intensively from less developed countries as vice versa.

To evaluate the causal impact of automation on trade and multinational activity, we exploit supply-side advances in the capabilities of robots over time that made it technically feasible to automate specific tasks. We follow Webb (2019) and use the text from patent titles to identify the tasks that robot-related patents substitute for, mapped to the occupations frequently conducting those tasks. This allows us to construct time-varying measures of exposure to automation with robots. Using this instrument, we show that for the full sample, starting to use robots in Spain caused firms to increase their value of imports from, and the number of affiliates in, less developed countries. Robot use also had a positive impact on the extensive margin of trade and multinational activity, leading firms to start importing from or start opening affiliates in less developed countries. For multinational activity, the expansion appears to be more directed at horizontal, rather than vertical, FDI, however, with robot use increasing the probability that the primary activity of the main affiliate is marketing or distribution of products manufactured in Spain, or producing similar products to those manufactured in Spain, as opposed to assembly or adaptation of inputs supplied by the firm in Spain.

We find that the sequencing of automation and offshoring matters. When we focus on firms that were importing intensively from less developed countries before they starting to use robots, we find that robot use had no effect on the value of imports but decreased the share of imports sourced from developing countries, suggesting that automation does, to some degree, shift economic activity away from developing countries, but only for firms that had offshored production first. By contrast, for firms that started using robots before importing intensively from less developed countries, robot use has a positive impact on the probability that they start to import from less developed countries and the value of their imports. The effect for the latter group dominates, such that the aggregate impact of robots on imports from less developed countries is positive.

In terms of the consequences for firm operations within Spain, we find that robot use also caused firms to substantially expand production and increase both

labour productivity and TFP, with some evidence of a decrease in the labour share. There is a weak, positive, impact of robot use on employment, but this is not robust to all specifications. Firms that start automating with robots are more likely to report that they have changed their regular workers, and experience a strong increase in the number of engineers and graduates employed and a weak increase in the number of production workers.

We demonstrate that these key results hold after controlling for global supply side shocks that have made offshoring cheaper and easier, import tariff changes, excluding the period of the Global Financial Crisis and excluding the automotive sector. For comparison, we also construct an additional instrument for robots using the capabilities of industrial robots as identified from robot sales data from the International Federation of Robotics (IFR), in a similar way to Graetz and Michaels (2018). We find similar results using this instrument for all outcomes except for the labour share. Finally, we also explore the impact of the intensive margin of automation, finding that increases in the stock of industrial machinery within firms that use robots also had similar impacts to the extensive margin of starting to use robots.

We construct a model that is able to provide a framework for understanding these findings by incorporating two new aspects into the canonical task-based automation model of Acemoglu and Restrepo (2018). First, we additionally incorporate Melitz (2003) style firm heterogeneity into the the task based framework and second, we allow for the additional option to offshore tasks as well as automate them, following from a variant in Acemoglu and Autor (2011) that incorporates offshoring in addition to automation. In a similar vain to Bustos (2011) and Antràs and Yeaple (2014) we model both technology adoption and offshoring as involving an up-front fixed cost, which provides the basis for why sequencing might matter, but offering a potential marginal cost reduction. An exogenous improvement in the share of tasks that can be automated generates a displacement effect, whereby labour at home or in an offshore location are replaced by machines, and a productivity effect, whereby automation increases productivity and allows firms to expand and demand more labour at home or abroad.

This research adds a number of new contributions to the existing, nascent literature on how automation in high income countries affects trade and FDI to less developed ones. There are now a few papers using industry-level robot sales data to study how robot adoption affects country-industry trade flows and Foreign Direct Investment (FDI). Artuc et al. (2018) and Hallward-Driemeier and Nayyar (2019) both show positive impacts of robot intensity in high-income countries on imports sourced from, or FDI growth to, less developed countries, respectively. We contribute by studying this same relationship at the firm level. Our results generally seem to reiterate this finding that the net effect of robot use on imports from less developed countries is positive. We also provide the contribution of offering empirical evidence and a modelling framework to explain the firm dynamics happening in the background: firstly, automation induces scale effects, allowing firms to expand and demand more from less developed countries. Secondly, there are a subset of firms that switch into importing from, or opening affiliates in, less developed countries as a consequence of automation.

Pedemonte et al. (2019), Artuc et al. (2019) and Faber (2018) on the other hand, focus only on trade between the US and Mexico, finding negative effects of automation. Using data on employment or exports in Mexican local labour markets, they show that increased robot penetration in the US reduces exports from Mexico to the US or employment in Maxico. Our research offers some hints as to why these results could hold, despite the other studies finding a positive relationship. Firstly, in our paper the impact of automation differed depending on whether firms first automated or first offshored. If the sample comprised of more firms that offshored first, we might expect a decrease in the share of imports from less developed countries for the full sample. Secondly, our modelling framework suggests that firms that don't automate could lose market share to those that do and, in our sample it holds that firms that don't automate, on average, have a higher import share from developing countries. This implies that a large enough market stealing effect could also result in a shift in the composition of imports away from developing countries.

Our paper also contributes to a small but growing literature studying automation using firm level data and evaluating the impact of automation on firm dynam-

ics. Our findings are consistent with the emerging work by Acemoglu et al. (2020) and Aghion et al. (2020) that automation induces within-firm scale and productivity effects. We additionally show that automation leads firms to change their regular workers, increase the wage bill and hire more engineers and graduates. Our results also suggest that studying the impact of automation on domestic outcomes in isolation of trade patterns could lead to incomplete or incorrect conclusions: in our analysis we provide evidence that automation does not only affect domestic outcomes, but the impacts of automation are also passed on to affiliates and trade partners.

This paper is also related to a wider literature studying the effects of robots on labour markets (Acemoglu and Restrepo, 2018, 2019a; Bessen et al., 2019; Graetz and Michaels, 2018). Finally, it also contributes to a growing literature studying how different technologies affect trade and global supply chains (Fort, 2017; Steinwender, 2018; Antràs, 2020; Baldwin and Forslid, 2020) and offers a new angle on how labour replacing technologies shape global trade. Finally it contributes to an expanding literature on superstar firms and on the reasons for the decline in the labour share (Autor et al., 2020; Song et al., 2019).

The rest of this paper proceeds as follows. In Section 1.2, we outline the model and discuss the potential relationship between automation and offshoring. In Section 1.3 we describe the data used. In Section 1.4 we explore the characteristics of firms that automate and the sequencing of automation, importing and multinational activity. Section 1.5 describes the empirical strategy. Section 1.6 provides the results, Section 1.7 provides robustness checks and Section 1.8 concludes.

1.2 Automation and offshoring with heterogeneous firms

This model incorporates two new aspects into the task based framework for studying automation of Acemoglu and Restrepo (2018). First, given we are studying automation at the firm level, we include Melitz (2003) style firm heterogeneity. Second, we include the additional option for firms to offshore, as well as automate, tasks. We take inspiration from Bustos (2011) and Antràs and Yeaple (2014), who model technology adoption and offshoring, respectively, in a Melitz framework.

1.2.1 Model setup

There is one monopolistically competitive industry where firms produce differentiated products under increasing returns to scale. Firms are heterogeneous in productivity as in Melitz (2003). Each firm produces a single variety ω and there is free entry. Firms heterogeneity is reflected by differing 'baseline' marginal costs of production $\varphi(\omega)$. To enter into production, firms pay a fixed entry cost of f_e units of labour. They then draw their productivity from a known Pareto cumulative distribution function $G(\varphi) = 1 - \varphi^{-k}$ with $k > 1$. After observing their productivity firms decide whether to exit or to produce. Following Acemoglu and Restrepo (2018) we assume that production is characterised by combining a unit measure of tasks according to a constant elasticity of substitution aggregator. The production of variety ω involves performing tasks $x \in [0, 1]$. The output of firm ω is then:

$$q(\omega) = \varphi(\omega)\left(\int_0^1 q(\omega, x)^{\frac{\sigma-1}{\sigma}} dx\right)^{\frac{\sigma}{\sigma-1}} \tag{1.1}$$

where $q(\omega, x)$ is the output of task x for firm ω and σ is the elasticity of substitution between tasks. We depart from Acemoglu and Restrepo (2018) but build upon a variant of Acemoglu and Autor (2011) by assuming that tasks can be performed by either labour at home $l_H(x)$, offshore labour $l_O(x)$ or automated machines $m(x)$.[3] Following Acemoglu and Restrepo (2018), we assume that labour in the home country, labour in the offshore country and machines are perfectly substitutable factors of production. Assume the tasks are ordered by the productivity of home labour in completing them, $\gamma_{LH}(x)$.

1.2.2 Production with only offshoring

We assume that only a subset of tasks $x \in [0, I^O]$ are possible to offshore because there are certain activities that cannot be performed at a distance and these are the tasks that home labour is the most productive in performing. For example, these could be activities that would typically be conducted at the firm headquarters in

[3]Machines could be used in either the home or offshore country. In our empirical analysis we can only observe use of robots in the firm in Spain, so focus on the use of machines at home. In general, the use of industrial robots per capita has been shown to be greater in high income countries, where robot manufacturing is generally concentrated and labour costs higher (IFR, 2019).

a high-income country. The tasks that are possible to offshore can be performed by labour at home or offshore labour, but performing a task with offshore labour involves an iceberg transport cost τ and performing any tasks with offshore labour involves paying a one-off upfront fixed cost f_O. Before automation becomes technologically feasible, the output of one task x is then:

$$q(\omega, x) = \begin{cases} \mathbb{1}[H = 1]\gamma_{LH}(x) l_H(\omega, x) + \mathbb{1}[O = 1]\frac{\gamma_{LO}(x) l_O(\omega, x)}{\tau} & \text{if } x \in [0, I^O] \\ \gamma_{LH}(x) l_H(\omega, x) & \text{if } x \in [I^O, 1] \end{cases} \tag{1.2}$$

where $\gamma_{LH}(x)$ is the productivity of home labour in task x, assumed to be increasing in x, $\gamma_{LO}(x)$ is the productivity of offshore labour in task x, $\mathbb{1}[H = 1]$ indicates that the firm chooses to conduct the task at home and $\mathbb{1}[O = 1]$ indicates that the firm chooses to conduct the task offshore. Analogously to in Acemoglu and Restrepo (2018), we assume that $\gamma_{LH}(x)/\gamma_{LO}(x)$ is increasing in x so labour has a comparative advantage relative to offshore labour in higher-indexed tasks.

Combining these equations, defining γ^{τ}_{LO} as the productivity of offshore labour adjusted to take into account iceberg costs and denoting $\Gamma_F^{(a,b)} = \left(\int_a^b \gamma_F(x)^{\sigma-1} dx \right)$, where F is either L_O, L_H or M, and $F^{(a,b)}(\omega) = \int_a^b f(\omega, x) dx$, where f is either l_H, l_O or m, we can write the firm's output if they choose to offshore $q^O(w)$ or if they do not offshore $q(w)$ as:

$$q^O(\omega) = \varphi(\omega) \left(\Gamma_{LO}^{(0,I^O)} L_O^{(0,I^O)}(\omega)^{\frac{\sigma-1}{\sigma}} + \Gamma_{LH}^{(I^O,1)} L_H^{(I^O,1)}(\omega)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{1.3}$$

and

$$q(\omega) = \varphi(\omega) \left[\Gamma_{LH}^{(0,1)} L_H^{(0,1)}(\omega)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{1.4}$$

1.2.3 Preferences and profit maximisation

Preferences across varieties have the standard CES form, with an elasticity of substitution $\sigma = 1/(1-\rho) > 1$. These preferences lead to demand function $q(\omega) = EP^{\sigma-1}[p(\omega)]^{-\sigma}$ for every variety ω , where $p(\omega)$ is the price of each variety, $P = [\int_0^M p(\omega)^{1-\sigma} d\omega]^{\frac{1}{1-\sigma}}$ is the price index of the industry, M is the measure of existing varieties and E is the aggregate level of spending in the country.

The profit maximising price is a constant markup over marginal costs. We consider the scenario where the unit cost of offshore labour is lower than the unit cost of home labour for tasks that are feasible to offshore. This implies that for the marginal task I^O we have that $\frac{\tau w_O}{\gamma_{LO}(I^O)} < \frac{w_H}{\gamma_{LH}(I^O)}$. In this scenario offshoring involves a marginal cost saving relative to using only home labour. In the absence of fixed costs, all firms would therefore offshore all tasks that are feasible to offshore and produce the others at home. In the presence of fixed costs, firms that offshore will offshore all of the tasks that are feasible to offshore. Then the marginal cost if the firm offshores, MC^O or doesn't offshore, MC is:

$$MC^O(\omega) = \frac{1}{\varphi(\omega)} \left(\Gamma_{LO}^{(0,I^O)} w_O^{1-\sigma} + \Gamma_{LH}^{(I^O,1)} w_H^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{1.5}$$

$$MC(\omega) = \frac{1}{\varphi(\omega)} \left(\Gamma_{H}^{(0,1)} w_H^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{1.6}$$

Taking the second part of the right hand side of 1.6, $\left(\Gamma_{H}^{(0,1)} w_H^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ as the numeraire and equating this to 1, we can then write the marginal cost functions as:

$$MC^O(\omega) = \frac{\alpha}{\varphi(\omega)} \qquad \text{and} \qquad MC(\omega) = \frac{1}{\varphi(\omega)} \tag{1.7}$$

where $\alpha < 1$ given our assumption that offshoring involves a marginal cost saving relative to using home labour.

Productivity cutoff for offshoring

If the firm offshores then they have profit function $\pi^O(\varphi(\omega))$, otherwise they have profit function $\pi(\varphi(\omega))$ as follows:

$$\pi^O(\varphi(\omega)) = (1-\rho) E P^{\frac{\rho}{1-\rho}} \left[\frac{1}{\rho} MC^O \right]^{\frac{-\rho}{1-\rho}} - F_O - F_e \tag{1.8}$$

$$\pi(\varphi(\omega)) = (1-\rho) E P^{\frac{\rho}{1-\rho}} \left[\frac{1}{\rho} MC \right]^{\frac{-\rho}{1-\rho}} - F_e \tag{1.9}$$

And so they will offshore if $\pi^O(\varphi(\omega)) > \pi(\varphi(\omega))$. There is a productivity cutoff φ^* associated with offshoring and firms sort into two groups: those which offshore

and those which don't. Then denoting $\Omega(\varphi(\omega)) = (1-\rho)EP^{\frac{\rho}{1-\rho}}\left[\frac{1}{\rho}\frac{1}{\varphi(\omega)}\right]^{\frac{-\rho}{1-\rho}}$ we have that firms will offshore when:

$$\pi^O(\varphi(\omega)) - \pi(\varphi(\omega)) > 0 \Leftrightarrow \Omega(\varphi(\omega))(\alpha^{\frac{-\rho}{1-\rho}} - 1) > F_O \tag{1.10}$$

and φ^* is the productivity corresponding to the firm where:

$$\Omega(\varphi^*(\alpha^{\frac{-\rho}{1-\rho}} - 1) = F_O \tag{1.11}$$

1.2.4 Introducing automation

We next assume that over time there is a subset of tasks $x \in [0, I^M]$ that become technically feasible to automate using existing technologies, such as robots, CNC machines and FMS. We begin by assuming that this subset of tasks is more limited than the subset that can be offshored, because certain tasks are not yet technically feasible to automate, for example, activities involving substantial manual dexterity or handiwork. That is, we assume $I^M < I^O$. Tasks $x \in [0, I^M]$ can be performed by either labour at home, offshore labour or machines. Performing any tasks with machines also involves paying a one-off fixed upfront cost f_M. The output of one task x is then:

$$q(\omega, x) = \begin{cases} \mathbb{1}[H=1]\gamma_{LH}(x)l_H(\omega,x) + \mathbb{1}[O=1]\frac{\gamma_{LO}(x)l_O(\omega,x)}{\tau} + \mathbb{1}[M=1]\gamma_M(x)m(\omega,x) & \text{if } x \in [0, I^M] \\ \mathbb{1}[H=1]\gamma_{LH}(x)l_H(\omega,x) + \mathbb{1}[O=1]\frac{\gamma_{LO}(x)l_O(\omega,x)}{\tau} & \text{if } x \in [I^M, I^O] \\ \gamma_{LH}(x)l_H(\omega,x) & \text{if } x \in [I^O, 1] \end{cases} \tag{1.12}$$

where $\gamma_M(x)$ is the productivity of machines in task x and $\mathbb{1}[M=1]$ is an indicator function denoting that the firm chooses to automate that task. We assume that $\gamma_{LH}(x)/\gamma_M(x)$ and $\gamma_{LO}(x)/\gamma_M(x)$ are increasing in x so labour at home and abroad has a comparative advantage relative to machines in higher-indexed tasks. We continue to assume that labour in the home country, labour in the offshore country and machines are perfect substitutes. Firms now face two additional prospective output functions for if they choose to automate as well as offshore q^{OM}, or if they choose only to automate q^M:

$$\mathrm{q}^{OM}(\omega) = \varphi(\omega)\left(\Gamma_M^{(0,I^M)} M^{(0,I^M)}(\omega)^{\frac{\sigma-1}{\sigma}} + \Gamma_{LO}^{(I^M,I^O)} L_O^{(I^M,I^O)}(\omega)^{\frac{\sigma-1}{\sigma}} + \Gamma_{LH}^{(I^O,1)} L_H^{(I^O,1)}(\omega)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \quad (1.13)$$

$$q^M(\omega) = \varphi(\omega)\left(\Gamma_M^{(0,I^M)} M^{(0,I^M)}(\omega)^{\frac{\sigma-1}{\sigma}} + \Gamma_{LH}^{(I^M,1)} L_H^{(I^M,1)}(\omega)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \quad (1.14)$$

We assume that the unit cost of machines is lower than the unit cost of offshore labour and the unit cost of home labour. This implies that for the marginal task I^M we have that $\frac{r}{\gamma_M(I^M)} < \frac{\tau w_O}{\gamma_{LO}(I^M)}$. In this scenario offshoring involves a marginal cost saving relative to using only home labour and automation involves a marginal cost saving relative to offshoring. In the absence of fixed costs, all firms would therefore automate all of the tasks that are technically feasible to automate, offshore the remainder that are feasible to offshore and produce the others at home. In the presence of fixed costs, firms that choose to automate will automate all of the tasks that are feasible to automate and firms that offshore will offshore all of the tasks that are feasible to offshore.

Profit maximisation

As before, the profit maximising price is a constant markup over marginal costs. The firms that have not chosen to automate or offshore therefore charge the highest price, while firms that both automate and offshore charge the lowest price. The marginal costs if the firm offshores and automates, MC^{OM} or only automates, MC^M are:

$$MC^{OM}(\omega) = \frac{1}{\varphi(\omega)}\left(\Gamma_M^{(0,I^M)} r^{1-\sigma} + \Gamma_{LO}^{(I^M,I^O)} w_O^{1-\sigma} + \Gamma_{LH}^{(I^O,1)} w_H^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \quad (1.15)$$

$$MC^M(\omega) = \frac{1}{\varphi(\omega)}\left(\Gamma_M^{(0,I^M)} r^{1-\sigma} + \Gamma_{LH}^{(I^M,1)} w_H^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \quad (1.16)$$

Once again taking the marginal cost without automation or offshoring as the numeraire we can write these marginal cost functions as:

$$MC^{OM}(\omega) = \frac{\beta}{\varphi(\omega)} \qquad \text{and} \qquad MC^M(\omega) = \frac{\delta}{\varphi(\omega)} \quad (1.17)$$

Our assumption that automation involves a marginal cost reduction relative to offshoring implies that $\beta < \alpha < 1$, while we also have that $\beta < \delta$ and $\delta < 1$.

The rank of α relative to δ, on the other hand, depends on the task subset that is feasible to automate, I^M relative to the subset feasible to offshore, I^O and the rental rate r relative to the cost of offshore labour w_O. To make the decision of whether to conduct tasks using only home labour, offshore labour, machines, or both offshore labour and machines, firms compare the profits under each option. Their profit functions if they automate, $\pi^M(\varphi(\omega))$, or if they offshore and automate, $\pi^{OM}(\varphi(\omega))$ are as follows:

$$\pi^M(\varphi(\omega)) = (1-\rho)EP^{\frac{\rho}{1-\rho}}\left[\frac{1}{\rho}MC^M\right]^{\frac{-\rho}{1-\rho}} - F_M - F_e \tag{1.18}$$

$$\pi^{OM}(\varphi(\omega)) = (1-\rho)EP^{\frac{\rho}{1-\rho}}\left[\frac{1}{\rho}MC^{OM}\right]^{\frac{-\rho}{1-\rho}} - F_M - F_O - F_e \tag{1.19}$$

There are now two additional productivity cutoffs associated with automating φ^{M*} and both automating and offshoring φ^{OM*}. Firms sort into four different groups depending on their productivity: the lowest productivity firms produce but do not automate or offshore, those with medium productivity either automate or offshore, and the highest productivity firms both automate and offshore. The benefit of automating and offshoring is firms earn higher revenues, because consumer demand is elastic ($\sigma > 1$), but firms must pay the higher fixed cost. The benefits of automating or offshoring are also increasing in firm productivity.

1.2.5 Advances in automation

This model allows for three types of technological progress in automation:

1. The extensive margin of automation: an increase in I^M.
2. The intensive margin of automation: an increase in $\gamma_M(x)$ or a decrease in r.
3. The fixed cost of automation: a decrease in f_M.

Here we focus on the extensive margin, as characterised by the invention of new technologies capable of automating a wider variety of tasks.

Firm level implications

At the firm level, holding I_O fixed, an increase in I_M has the following effect:

- **Firms that are already automating** will conduct more tasks using machines, reducing demand for offshore labour or labour at home. However, the improvement will also lower marginal costs, meaning that the firms that automate can reduce prices, increase revenues and expand, increasing demand for offshore labour or home labour in the tasks that can't be automated. Automation at the extensive margin in our model therefore has the same two counteracting effects as in Acemoglu and Restrepo (2018) that affect both demand for offshoring and for home labour within the firm: the negative displacement effect and the positive productivity effect.

- **Firms that are induced to automate** because of the further marginal cost reduction of being able to use machines for more tasks will experience a larger displacement effect as they switch a whole subset of tasks $x \in [0, I^M]$ away from either home labour if they hadn't offshored or offshore labour if they had, towards machines, and also a positive productivity effect.

Industry level implications:

An expansion in the task subset that can be automated reduces the productivity cutoff associated with automating and so raises industry level automation. In turn this raises the survival cutoff, meaning that firms that cannot automate are forced to exit and surviving non-automating firms reduce their output, employment and offshoring. In the world with more automation, the expected productivity level of surviving firms is therefore higher than in the world without automation and the per period expected profits of surviving firms are also higher.

1.3 Data

1.3.1 Spanish firms data

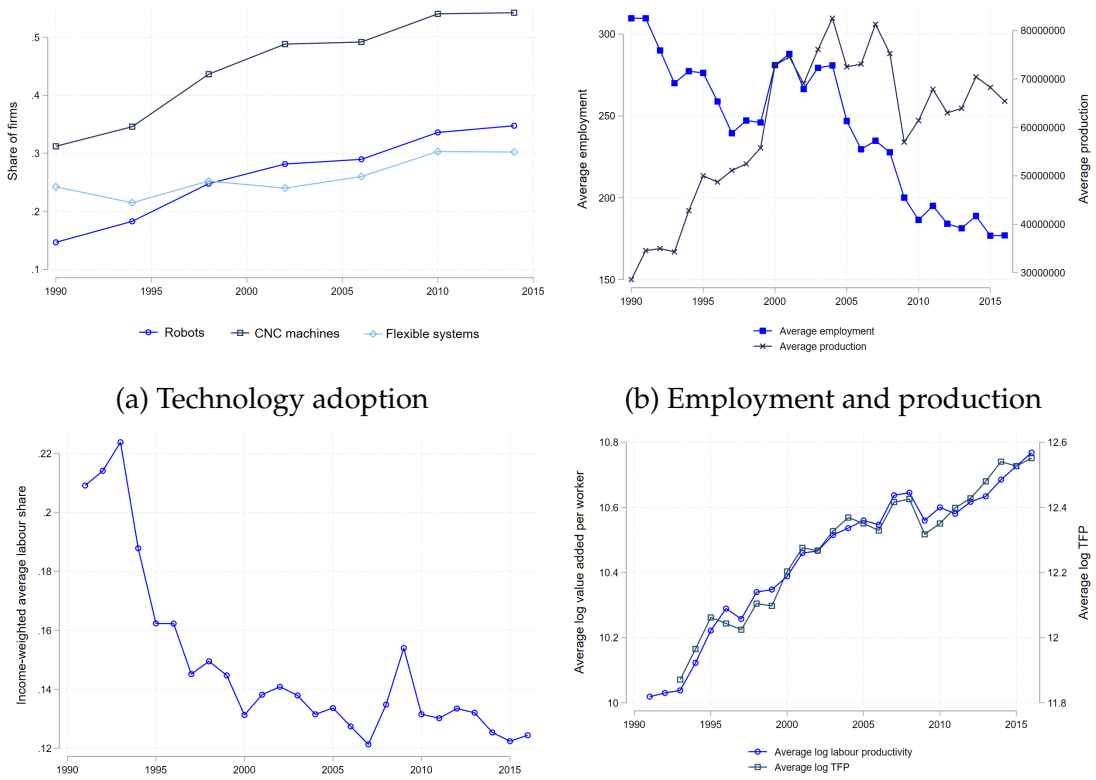
The primary data source used in this paper is the 'Encuesta Sobre Estrategias Empresariales' (ESEE), an annual survey of Spanish manufacturing firms carried out

by the SEPI Foundation in conjunction with the Spanish Ministry of Industry between 1990 and 2016.[4] This dataset is relatively unique in providing a rich set of variables on firm technology use and information on trade and multinational activity. The survey generates an unbalanced panel of 5,840 legal entities for the years 1990-2016. The SEPI Foundation applies a complex random sampling procedure, sending out survey questionnaires to all firms with more than 200 employees, and to a subset of firms with 200 or less but more than 10 employees. This subset is selected according to a stratified sampling scheme that guarantees that they can establish representativeness of the data for different industries and the manufacturing sector as a whole. The SEPI Foundation preserves these sample properties over time by controlling for the dynamics in the panel due to market entry and exit. In total there are 49,237 firm-year observations, with an average of 1824 firms in each year and a median duration in the sample of six years.

Automation related variables: The dataset provides dummy variables on firm use of three distinct technologies related to automation: robots, CNC machines and FMS. These variables are reported for four year periods. The dataset also asks about firm use of Computer Aided Design/Computer Aided Manufacturing (CAD/CAM), but in our analysis we don't include CAD/CAM as CAD/CAM is a process and software that is typically combined with CNC machines, where CNC machines are the machinery that conducts the task. A description of these technologies is as follows:

1. **Robots:** Robots used in the manufacturing sector typically perform industrial applications including welding, material handling, assembly, palletizing, and painting. The IFR defines industrial robots by the ISO 8373 definition: 'An automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use in industrial automation applications' (IFR, 2019). The ESEE does not specify the type of robots used. However, given the survey covers manufacturing firms and up to 2016, when collaborative robots

[4] https://www.fundacionsepi.es/investigacion/esee/en/spresentacion.asp



(a) Technology adoption

(b) Employment and production

(c) Labour share

(d) Productivity

Figure 1.1: Key domestic outcomes over time

Notes: Panel (a) displays the share of firms in the sample that report using the three technologies. Panel (b) displays the mean value of production of goods or services and the mean employment across all firms. Panel (c) displays the income-weighted average labour share across all firms. Panel (d) displays the mean logged value added per worker and logged TFP across all firms.

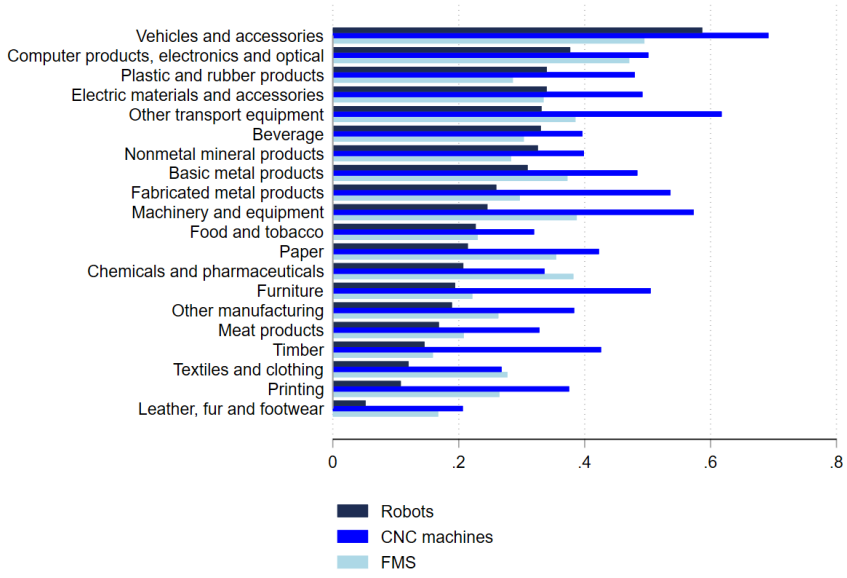


Figure 1.2: Share of firms using each technology, by industry

Notes: This figure displays the share of firms using each of the three technologies across all time periods, by industry.

were less widely commercially available, we assume these are likely to be primarily industrial robots.

2. **FMS:** An FMS combines CNC machines, industrial robots, and other types of automation into one fully automated system. A FMS would typically produce similar products and parts but still maintain the flexibility to change parts or processes.

3. **CNC Machines:** CNC machines use computers to store, calculate, and execute operations that are usually performed by hand. A common example of a CNC machine is a CNC mill, which uses computers to analyze, cut, and mill each piece of material.

Figure 1.1(a) displays the share of firms using these technologies over time and Figure 1.2 displays the share using these technologies on average across all years by industry.[5] There are 20 industries in the ESEE data using the Spanish NACE-

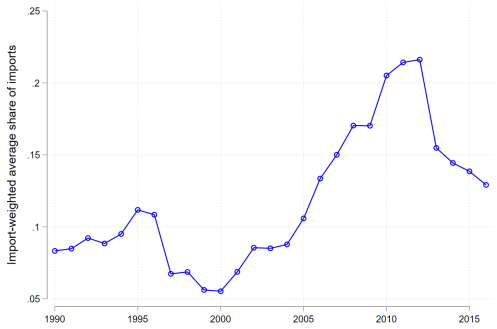
[5]In Appendix 1.A, Figure 1.8 also displays technology adoption by region.

CLIO classification system, which is based on the NACE.[6] The share of firms using robots and CNC machines has approximately doubled over our sample period, while the share using FMS has risen more moderately. The most automated industries are Vehicles and Accessories, Computer Products, Electronics and Optical and Plastic and Rubber Products, while the least automated are Leather, Fur and Footwear, Printing and Textiles and Clothing. In Table 1.16 in Appendix 1.B we also display the correlation between use of these three technologies, which is highest for robots and CNC Machines at 0.34 and lowest for FMS and CNC Machines at 0.18.

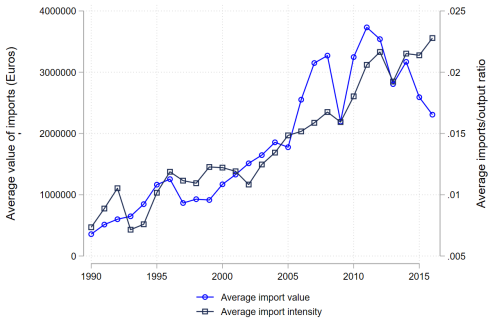
Additionally the ESEE data provides annual investment by type, with one of the categories the 'purchase of technical facilities, machinery and tools'. We use this, along with the firm level depreciation rate, to construct the firm level stock of industrial machinery and equipment, discussed later in the paper.

Trade and multinational activity variables: The key trade-related variables we use are the value of imports and breakdown of imports by region. Firms provide exports and imports broken down into region groups: the EU, OECD countries excluding the EU, Latin America and the Rest of World. We combine the sum of Latin-America and Rest of World to form a broad measure of trade with 'less developed countries'. We list the countries in these groups in Appendix 1.C. In terms of total import value cumulatively between 1990 and 2016, Spain's largest import partners in this less developed group were China, with 5.5% of imports, Algeria, with 1.8%, Nigeria with 1.6%, Russia with 1.5%, Saudi Arabia with 1.2%, Brazil with 1.1%, Morocco with 1%, Libya with 0.8%, India with 0.8% and Indonesia with 0.6%, based on UNCOMTRADE import data. Membership of the EU and OECD changed slightly over time and we document these changes in Table 1.23 in Appendix 1.C. The group of countries that switch into the less developed category during our sample period account for a relatively small amount of Spain's imports, at around 6%. We include industry group-year fixed effects in our empirical

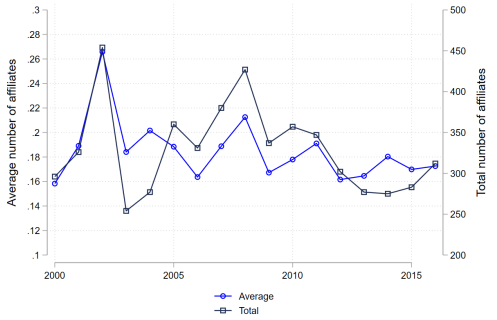
[6]These industries are meat related products; food and tobacco; beverage; textiles and clothing; leather, fur, and footwear; timber; paper; printing and publishing; chemicals; plastic and rubber products; nonmetal mineral products; basic metal products; fabricated metal products; industrial and agricultural equipment; office machinery, data processing, precision instruments and similar; electric materials and accessories; vehicles and accessories; other transportation materials; furniture; other manufacturing.



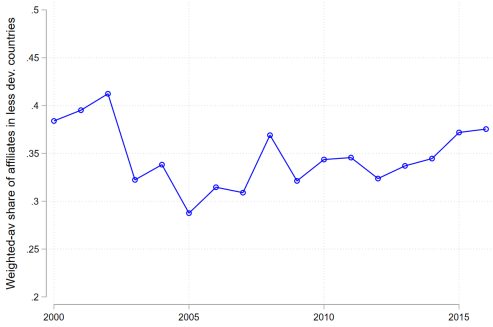
(a) Import share from less dev. countries



(b) Value of imports & import intensity



(c) Affiliates in less dev. countries



(d) Share of affiliates in less dev. countries

Figure 1.3: Key international outcomes over time

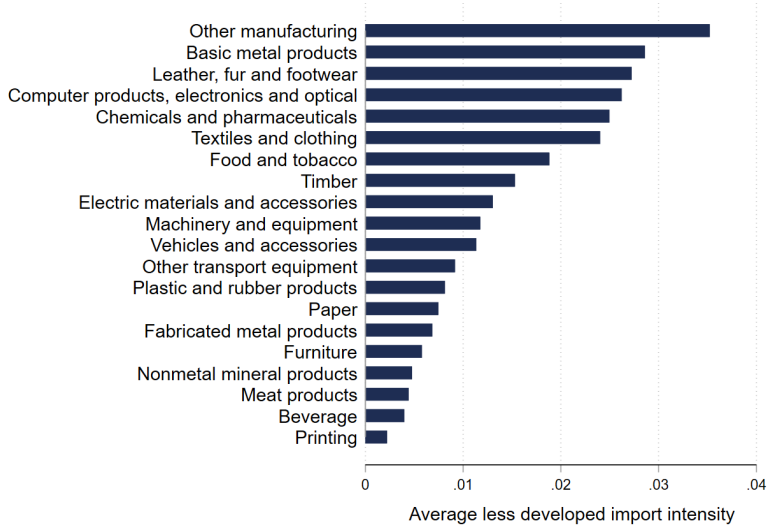
Notes: Panel (a) displays the import-weighted average share of firm imports that originate from less developed countries, averaged across all firms. Panel (b) displays the mean value of imports from less developed countries and the mean import intensity from less developed countries, defined as the value of imports from less developed countries scaled by firm output. Panel (c) displays the mean number of affiliates in less developed countries and the total number, across all firms. Panel (d) displays the weighted-average share of affiliates in less developed countries, across all firms.

specifications, which should absorb these changes to categorisation that we expect would affect all firms equally, or at least all firms in the same industry groups equally.

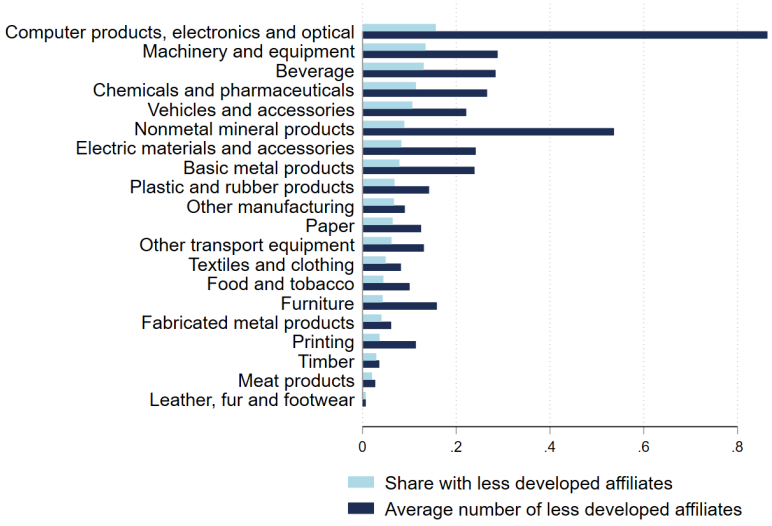
Our primary outcome variables are a dummy variable for whether a firm imports from this group of less developed countries, the share of imports from this group and the intensity of imports from this group, measured as imports relative to output. Additionally, firms provide their number of affiliates by these same region groups, the country of the main affiliate, employment at the main affiliate and information on the activities of affiliates. They ask whether the main activity of the affiliated company consists only of commercialisation or distribution, whether the main affiliated company manufactures similar products to those manufactured in Spain by its company or whether the main affiliated company carries out adaptation and/or assembly activities of parts provided by the company in Spain. In our baseline analysis we focus on two key variables: the number of affiliates in less developed countries and the share of affiliates in less developed countries. We then explore changes in the employment and main activities as a secondary exercise.

Our model framed our analysis in terms of 'offshoring' of tasks. There are a range of definitions of offshoring that have been used in the literature. Blinder and Krueger (2013) define offshoring as 'The movement of home-country jobs to another country'. Hummels et al. (2014) define it as 'The import of intermediate inputs that could have been produced within the same firm'. Feenstra and Hanson (1999) distinguish between narrow and broad offshoring, where the former considers imported intermediates in a given industry from the same industry only, while the latter considers imported intermediates from all industries. Our import variable is not only intermediate imports, but all firm imports, so is slightly broader than these definitions. However, from the year 2006 onwards, the ESEE data starts to ask firms about the share of their imports that are imported inputs. The median response is that 75% of imports are imported inputs, suggesting that our import variable is a close proxy measure of 'broad' offshoring.

Figure 1.3 displays the key trade and multinational activity variables over time. Between 1990 and 2016 the mean value of imports from less developed countries and the intensity of imports from less developed countries, have increased ap-



(a) Import intensity



(b) Affiliates

Figure 1.4: Mean less developed import intensity, share with affiliates & no. of affiliates, by industry

Notes: Panel (a) displays the mean import intensity from less developed countries across all years, by industry, where import intensity is defined as the value of imports from less developed countries scaled by output. Panel (b) displays the share of firms with at least one affiliate in a less developed country and the mean number of less developed affiliates, across all years, by industry.

proximately six-fold.[7] The import-weighted share of imports from less developed countries approximately tripled between 1990 and 2010 but has then fallen. There is a less clear pattern for affiliates in less developed countries. Figure 1.4(a) displays the import intensity from less developed countries by industry. The industries with the highest import intensity are Other Manufacturing, Basic Metal Products and Leather, Fur and Footwear. Figure 1.4(b) displays the share of firms with affiliates in less developed countries and the average number of affiliates in less developed countries, by industry. The three industries with the highest share of firms having affiliates in less developed countries are Computer products, electronics and optical, Machinery and equipment and Beverage.

Measuring TFP: Our main measure of TFP is constructed using the Levinsohn-Petrin method, which overcomes the endogeneity problem associated with estimating TFP using OLS by using intermediate inputs to control for unobserved productivity changes. We choose this method over Olley-Pakes, which uses investment as the control variable, because investment is often reported as zero (in our sample it is zero or missing in just over 14% of observations), while there is better data coverage for intermediate inputs (zero or missing for only 0.6% of observations). The ESEE data includes firm level changes in input and output prices. It asks firms about the annual change in the sales price of their products, the change in the purchased price of energy, raw materials and services. We use the sales price change to deflate the value of total income, which is the sum of firm production of goods and services and other income, the average of the change in energy and raw materials prices to deflate the value of intermediate purchases of goods inputs and the services price change to deflate the value of intermediate purchases of services inputs.

The ESEE data doesn't provide firm level price changes for capital, so we use the gross fixed capital formation price index for the manufacturing sector for Spain from the Capital Input Files from EUKLEMS.[8] We construct real value added as the real value of total income minus the real value of intermediate inputs of goods

[7] The same holds for the total value of imports from less developed countries across all firms included in the data.

[8] http://www.euklems.net/

and services and we use real value added as the dependent variable. We measure labour in terms of the average total employment at the firm, provided in the ESEE data. This is calculated as the sum of the following items: Full-time regular personnel, 1/2 of the part-time regular personnel and the average number of temporary workers. Given our data only covers manufacturing industries and the sample size too small to estimate TFP for individual sub-manufacturing sectors, we instead split the manufacturing industries into two broad groups and estimate TFP for each group.[9] As a robustness check, in Appendix 1.B we also repeat all results using TFP constructed by the Olley-Pakes method instead, showing that there is only a slight difference in magnitudes of coefficients between these two estimates.

Employment categories: In addition to providing total employment, the data also includes employment by a number of categories: the share of engineers and graduates, the share of non-graduates, the share of production workers and the share of clerical workers, where clerical workers is defined to include managerial staff and administrative staff.

Other variables: The main analysis also explores how automation affects total employment, output, which is defined as the total value of production of goods and services, value added per worker and the labour share of income.[10] Figure 1.1 displays the evolution of these variables over the sample period. Table 1.1 displays the summary statistics for the key variables used in the analysis.

1.3.2 Spanish employment data

We obtain data on the composition of employment in 1981 by 2 digit ISCO 68 occupation code within each industry and region from IPUMS International, which

[9]These groups are: meat related products; food and tobacco; beverage; textiles and clothing; leather, fur,and footwear; timber; paper; printing and publishing; chemicals; plastic and rubber products in one group and nonmetal mineral products; basic metal products; fabricated metal products; industrial and agricultural equipment; office machinery,data processing, precision instruments and similar; electric materials and accessories; vehicles and accessories; other transportation materials; furniture; other manufacturing in the other group

[10]Following consensus in the literature we define the labour share as total labour costs, divided by total income, both variables that are directly reported in the ESEE data.

TABLE 1.1. SUMMARY STATISTICS

	Mean	S.d.	Min	Max
Variables in levels				
Robot user	0.27	0.44	0	1
FMS user	0.28	0.45	0	1
CNC machine user	0.45	0.50	0	1
Robot user*machinery stock	6.8e+14	1.1e+17	0	2.32e+19
Patent exposure measure	33.8	20.1	5.00	118.4
IFR exposure measure	0.28	0.26	0	0.89
IFR robot stock (millions)	0.92	0.34	0.56	1.83
Share of imports from less dev countries	0.44	0.47	0	1
Imports from less dev countries (euro millions)	1.85	22.5	0	1626.5
Import intensity from less dev countries	0.014	0.055	0	1.54
Affiliates in less dev countries	0.18	1.55	0	100
Share of affiliates in less dev countries	0.043	0.18	0	1
Employment	239.6	767.8	1	24634
Output (euro millions)	59.7	292.1	5.4	8923.4
Labour share	0.29	0.35	0.0049	59.0
Transformed variables				
Robot user*log machinery stock	4.35	7.02	0	44.6
IHS(imports from less dev countries)	3.03	5.82	0	21.9
IHS(import intensity from less dev countries)	0.014	0.053	0	1.21
IHS(affiliates in less dev countries)	0.092	0.39	0	5.30
Log employment	4.20	1.50	0	10.1
Log output	15.8	2.03	8.59	22.9
Log labour productivity	10.4	0.72	1.50	15.2
Log TFP	12.3	1.24	1.21	17.9

Notes: This table shows the summary statistics for the key variables used in the main empirical analysis, pooled across all years.

obtains its data from the Spanish Census of Population and Housing 1981. This census covers a 5% sample of the population. We choose the 1981 sample because we are studying the impact of automation on offshoring and thus want the employment composition before offshoring began on a large scale. Ideally we would like to use the full sample of employment by industry, region and occupation, but this does not exist for Spain at a detailed enough level of disaggregation. The Census of Housing and Population is the most comprehensive data source on employment in these three dimensions.

1.3.3 Patent and occupation data

The construction of the robot exposure measures as in Webb (2019) uses Google Patents Public Data, provided by IFI CLAIMS Patent Services.[11] The fields used

[11]https://bigquery.cloud.google.com/table/patents-public-data:patents.publications201710

are the title, abstract, and CPC codes, as described further below. The O*NET database, produced by the US Department of Labor, is also used as the source of information on occupations. O*NET describes 964 occupations. A set of tasks is listed for each occupation, described in natural language. Each task is also given scores that indicate its importance and frequency in the occupation. These scores are used to weight tasks within occupations. The method for constructing patent-occupation text similarity is outlined in Section 1.5.

1.3.4 IFR robot sales data

We also make use of data from the IFR. The IFR measures global shipments of 'multipurpose manipulating industrial robots', based on the ISO definition provided above. The IFR data includes shipments delivered to each country by industry and application for the time period 1993-2016. Typical applications of industrial robots include assembling, dispensing, handling, processing (e.g., cutting), and welding, all of which are prevalent in manufacturing industries. We use this data to construct additional robot exposure measures, closely following the method of Graetz and Michaels (2018), discussed further in Section 1.5.

1.3.5 Offshoring instrument control

As a robustness check, we develop a control variable for exogenous supply-side factors that may have affected a Spanish firm's propensity to offshore production, that are unrelated to automation. We construct a variable similar to the 'World Export Supply' instrument used in Hummels et al. (2014) that aims to soak up supply shocks occurring within Spain's import partner destinations, for example, the rise of China as a manufacturing hub. To do so, we make use of the import share variables in the ESEE data that break down imports by the four regional groups discussed above. We use these shares in the firm's first reporting period to weight industry level exports from all countries in these region groups to four EU countries with the closest GDP per capita to Spain.[12]

We take this export data by ISIC Rev 3 industry from UNCOMTRADE, accessed via the WITS platform from the World Bank.[13] We use an ISIC Rev 3-NACE Rev

[12]These countries are Italy, Portugal, Slovenia and the Czech Republic.
[13]wits.worldbank.org

1 Crosswalk provided by Eurostat to map these to NACE Rev 1 industries.[14] The Spanish NACECLIO industry classification is based upon NACE Rev 1 so we can then easily map these to the ESEE industries. In order to ensure the measures reflect supply shocks in the industries that Spanish firms are likely to source their intermediate imports from, we further apportion the industry level exports to other industries according to the baseline purchase share specified in Spain's I-O table for imports obtained from the Instituto Nacional D'Estistica (INE). [15] Then for firm i in industry j at time t, the instrument for offshoring O_{ijt} is defined as:

$$O_{ijt} = \sum_{c} s_{ijc,base} \times \text{IO share}_{jk,base} \times EX_{ckt} \tag{1.20}$$

where $s_{ijc,base}$ is firm i in industry j's share of imports from region group c in their first reporting period, IO share$_{jk,base}$ is the IO share of sourcing industry k in purchasing industry j's total imports in 1995 and EX_{ckt} is the value of total exports from region group c in industry k to the four countries similar to Spain.

1.3.6 Import tariff controls

As a robustness check we also develop a control variable for changes to Spain's import tariffs that may have affected the cost of importing from less developed countries or the opening of affiliates in less developed countries. We follow the method used by Chen and Steinwender (2019), which exploits the fact that tariffs for Spain are set at the EU level and hence plausibly exogenous. We construct two industry-year level variables for MFN import tariffs on goods in the same industry ('same industry import tariffs') and MFN import tariffs on the imported inputs used by that industry ('imported input tariffs'). We take the data for MFN tariffs from TRAINS (provided by UNCTAD) accessed via the WITS platform. We take the simple average tariff in each ISIC Rev 3 product category in each year and use the ISIC Rev 3-NACE Rev 1 Crosswalk discussed above to map these to NACE Rev 1 industries. We then use UNCOMTRADE data from WITS on Spanish imports in 1990 by ICIC Rev 3 to create baseline-import weighted average tariffs for each NACE Rev 1 industry. For the imported input tariffs we further weight

[14] https://ec.europa.eu/eurostat/ramon/relations/index.cfm?TargetUrl=LST_REL

[15] https://www.ine.es/dyngs/INEbase/en/operacion.htm?c=Estadistica_C&cid=1254736177058&menu=enlaces&idp=1254735576581

TABLE 1.2. PREVALENCE & MARKET SHARE OF FIRMS THAT AUTOMATE

	1990	2002	2014	1990	2002	2014	1990	2002	2014
		Robots			CNC machines			FMS	
					Total number of firms				
Don't use	1856	1,638	1375	1494	1,177	956	1651	1,733	1450
Use	321	669	686	683	1,128	1105	518	574	611
% using	15%	29%	33%	31%	49%	54%	24%	25%	30%
					Total employment (millions)				
Don't use	0.29	0.18	0.10	0.24	0.13	0.09	0.29	0.21	0.12
Use	0.24	0.28	0.19	0.29	0.32	0.18	0.34	0.25	0.16
% using	46%	61%	64%	56%	71%	67%	54%	54%	57%
					Total sales (Euro billions)				
Don't use	28.1	37.1	36.1	29.0	31.9	26.1	19.2	43.4	36.3
Use	31.2	80.1	71.6	30.3	85.3	81.7	40.1	73.8	71.4
% using	53%	68%	66%	51%	73%	76%	68%	63%	66%
					Total exports (Euro billions)				
Don't use	4.8	10.4	16.3	4.1	7.8	11.5	2.8	12.4	16.3
Use	4.9	33.6	41.2	5.7	36.1	46.1	6.9	31.6	41.3
% using	51%	76%	72%	58%	82%	80%	71%	72%	72%
					Total imports (Euro billions)				
Don't use	3.47	6.74	8.79	2.9	5.4	4.1	1.95	8.47	6.54
Use	5.56	23.3	24.6	6.1	24.6	29.3	7.07	21.6	26.8
% using	62%	78%	74%	67%	82%	88%	78%	72%	80%

Notes: This table shows the total number of firms, the total number of employees, the sum of the value added, production and exports of all firms in the sample, broken down by whether or not they use robots, CNC machines and FMS, in a given time period.

these tariffs using the IO import table weights obtained from the INE, as described above.

1.4 Which firms automate?

We begin by documenting the characteristics of firms that use robots, FMS and CNC machines over time. Our modelling framework incorporates Melitz (2003) style heterogeneity and here we demonstrate why this is an important aspect for understanding the causes and consequences of automation. In terms of market share, manufacturing firms using our three technologies are in the minority, but make up the majority of employment, sales, value added and exports, as is shown in Table 1.2. In 1990 15 percent of the firms in the sample used robots, but accounted for 46 percent of employment, 53 percent of value added and output and 51 percent of exports. The proportion of firms using CNC machines is higher than all other technologies.

TABLE 1.3. TECHNOLOGY ADOPTION PREMIA

	1994	2002	2014	1994	2002	2014	1994	2002	2014
		Robots			CNC machines			Flexible systems	
				Panel (a): Firm characteristics					
Log total employment	1.35	1.37	1.06	0.93	0.85	0.62	1.29	1.09	0.91
Log value added per worker	0.19	0.29	0.25	0.13	0.19	0.15	0.29	0.23	0.25
Log TFP	0.74	0.81	0.63	0.50	0.48	0.34	0.77	0.64	0.57
Log output	1.70	1.79	1.40	1.18	1.06	0.86	1.68	1.42	1.27
Labour share	-0.039	-0.060	-0.078	-0.047	-0.036	-0.035	-0.051	-0.043	-0.047
				Panel (b): International Trade					
Exporter	0.28	0.27	0.17	0.18	0.17	0.12	0.22	0.19	0.14
Importer	0.27	0.27	0.18	0.19	0.16	0.11	0.26	0.17	0.16
IHS(exports)	5.40	5.54	3.73	3.29	3.22	2.52	4.22	3.93	3.14
IHS(imports)	4.97	5.15	3.70	3.25	2.96	2.23	4.65	3.38	3.31
IHS(imports from less developed)	1.57	1.70	2.53	1.03	1.15	1.76	1.50	1.85	2.03
Import share from less developed	-0.27	-0.26	-0.15	-0.18	-0.15	-0.086	-0.22	-0.15	-0.14
				Panel (c): Multinational activity					
Multinational	0	0.15	0.11	0	0.066	0.051	0	0.082	0.097
IHS(total number of affiliates)	0	0.24	0.21	0	0.11	0.12	0	0.18	0.16
IHS(number of less developed affiliates)	0	0.10	0.092	0	0.039	0.073	0	0.099	0.062

Notes: this table shows the coefficients from regressions of the dependent variables listed on the left hand side on a dummy variable for firm use of one of the three technologies in the specified year. All coefficients are significant to the 1 % level. These regressions also included industry and region fixed effects. We take the Inverse Hyperbolic Sine (IHS) transformation due to the presence of zeros. Standard errors were clustered at the industry-region level in all regressions.

1.4.1 The performance gap for firms that automate

We further document the performance gap for firms that use these technologies compared with firms that don't. We report these differences in Table 1.3 for three time periods (1994, 2002 and 2014). This table reports the technology adoption 'premia' estimated from a regression of the form:

$$X_i = \alpha + \beta \text{Technology}_i + \gamma \text{Industry}_i + \theta \text{Region}_i + \epsilon_i \tag{1.21}$$

where X_i is the outcome of interest, Technology_i is a dummy variable for firm use of one of the three technologies: robots, CNC or FMS, Industry_i is the firm's 2 digit industry and Region_1 is the Spanish autonomous community where the firm is headquartered.[16] Each premium shows the average difference between firms in the same industry and region using or not using that specific technology. Panel (a) of Table 1.3 reports these premia, where for all of the technologies in all of the

[16]The ESEE data includes 20 regions.

time periods these premia are statistically significant to the 1% level for all firm characteristics listed. Across all the technologies and time periods, firms using these technologies employ approximately twice the number of workers, have 22% higher labour productivity, 60% higher TFP, produce more than double and have a labour share that is 0.05 lower. However, the size and productivity premia tend to be decreasing over time, suggesting that barriers to automating might have been decreasing over time.

Panel (b) of Table 1.3 shows the difference in international orientation for firms using these technologies. In all time periods, firms using these technologies are on average 20% more likely to be exporters or importers, import and export nearly four times more to all destinations, import nearly twice as much from less developed countries but the share of their imports from less developed countries is around 17% lower. Panel (c) shows that firms that automate are, on average, also 6% more likely to be multinational corporations, have around 11% more affiliates on average and 5% more affiliates in less developed countries. The premia for the share of affiliates in less developed countries were not statistically significant.

1.4.2 Selection into automation

While it is clear that firms that automate have distinct characteristics, it is not clear if this is the impact of automation or a selection effect. We hence also consider the characteristics of firms that will select into automating in the future, prior to having started. We calculate the *ex ante* premia for firms that will use these technologies at any point in the future, but do not yet use them in the present, relative to firms in the present who never adopt these technologies in the sample period. Table 1.4 documents the coefficients for regressions of the form:

$$X_{it} = \alpha + \beta \text{FutureTechnology}_{it} + \gamma \text{Industry}_{it} + \theta \text{Region}_{it} + \epsilon_{it} \tag{1.22}$$

where X_{it} is the outcome of interest for firm i in period t, $\text{FutureTechnology}_{it}$ is a dummy variable for firm use at any point in the future recorded in the data of one of the technologies: robots, CNC machines or FMS, Industry_{it} is firm i's 2 digit industry at time t and Region_{it} is the Spanish autonomous community where firm i is headquartered at time t. The sample is limited to firms not currently using

TABLE 1.4. FUTURE TECHNOLOGY ADOPTION PREMIA

	1994	2002	2010	1994	2002	2010	1994	2002	2010
		Robots			CNC machines			Flexible systems	
				Panel (a): Firm characteristics					
Log total employment	1.15	0.85	0.92	0.86	0.76	0.69	1.06	0.88	0.96
Log value added per worker	0.26	0.22	0.17	0.27	0.20	0.22	0.27	0.36	0.25
Log TFP	0.65	0.55	0.59	0.53	0.53	0.56	0.65	0.62	0.61
Log output	1.49	1.17	1.20	1.24	1.13	1.03	1.40	1.19	1.29
Labour share	-0.057	-0.062	-0.073	-0.079	-0.068	-0.073	-0.063	-0.048	-0.064
				Panel (b): International trade					
Exporter	0.24	0.20	0.17	0.19	0.17	0.16	0.23	0.20	0.17
Importer	0.24	0.20	0.21	0.21	0.18	0.12	0.25	0.21	0.16
IHS(exports)	4.16	3.59	3.47	3.34	2.92	3.16	4.05	3.89	3.56
IHS(imports)	4.23	3.57	3.72	3.44	3.03	2.31	4.15	3.59	3.09
IHS(imports from less developed)	0.88	0.94	1.64	1.06	1.36	1.09	1.10	0.93	1.01
Import share from less developed	-0.23	-0.19	-0.20	-0.19	-0.16	-0.092	-0.24	-0.20	-0.16
				Panel (c): Multinational activity					
Multinational		0.083	0.092		0.066	0.080		0.14	0.12
IHS(total number of affiliates)		0.15	0.11		0.13	0.100		0.23	0.17

Notes: this table shows the coefficients from regressions of the dependent variables listed on the left hand side on a dummy variable for if a firm begins using one of the three technologies at any point in the future. The sample is limited to firms not using these technologies in the present. These regressions also controlled for industry and region dummy variables. All coefficients were significant to the 5% level. Standard errors were clustered at the industry-region level in all regressions.

the given technology. Each premium shows the average difference between firms in the same industry and region that start adopting the specific technology in the future compared to those that never adopt it. We limit the sample to only firms that remain in the sample for at least 5 years so that we compare firms where we have the information on whether they adopted these technologies in the future.

Table 1.4 displays all of the premia that are statistically significant to the 5% level, showing that the distinct characteristics of technology adopters also hold before adoption happens, except for the number of affiliates in less developed countries. These premia tend to be lower and less statistically significant than for actual adopters, suggesting that automation might also have a causal impact on these firm outcomes, or firms that adopt technology sooner rather than later also have distinct characteristics.

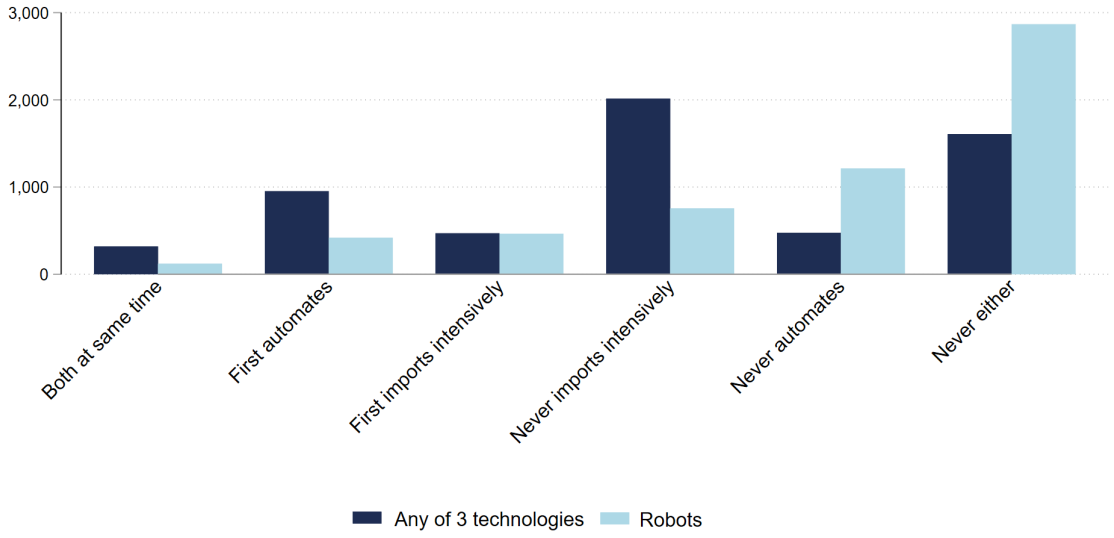


Figure 1.5: Sequencing of automation and importing intensively from less developed countries

Notes: This figure displays the number of firms by their sequencing of automation and importing intensively or opening affiliates in less developed countries. Importing intensively is defined as being in the top 20% of the sample in terms of import intensity from less developed countries.

1.4.3 Which comes first? Automation or importing/opening affiliates

We have shown that firms that automate, on average, import more and have more affiliates in less developed countries, although their import share from less developed countries is lower. But it is not yet clear which is causing the other. We hence further consider the timing of when firms automate and when they start importing intensively or opening up affiliates in less developed countries. We define importing intensively from less developed countries as being in the top 20th percentile in terms of the ratio of imports from less developed countries to output. We define automating as using either robots, CNC machines or flexible manufacturing systems.

The debate around reshoring typically assumes that firms first offshored production, then returned it home again. If automation is a cause of reshoring, then we might expect that firms would first report either importing intensively from, or having affiliates in, developing countries, then later report automating, followed

by a decline in import intensity or affiliates in less developed countries. Figure 1.5 shows that, in fact, far more firms started automating, using any one of these three technologies, before they started importing intensively from less developed countries. Only 36% of firms ever import intensively from less developed countries, while 64% automate at some point. The most common outcome is for firms to automate but never import intensively. Of those that do both, they are more than twice as likely to report automating before they import intensively. In terms of only those that automate only using robots, which is a smaller fraction of firms, only 30% of firms use robots at some point in the sample, and using robots is less common than importing intensively from less developed countries. For firms that do both, they are almost equally likely to be importing intensively before they automate as vice versa.[17]

Having an affiliate in a less developed country is far less common than importing intensively from less developed countries. Only 6% of our sample ever report having affiliates in less developed countries. Our variables on affiliates were only introduced after 2000 and so we cannot compare the timing of those that do both, but we show that it is very rare for firms in our sample to report having affiliates in developing countries but never automating. Of those that have affiliates in less developed countries, only 17% never automate using one of the three technologies.

1.4.4 Productivity of firms that automate, import and have affiliates

We assumed in our model section that firms select into automating and offshoring depending on their productivity, with the most productive firms both automating and offshoring and the least productive firms doing neither. We show in Figure 1.6 that this is the case in our data. We define automating and importing intensively as above. Firms that both automate and import intensively have a labour productivity distribution that is a shift right relative to firms that do one or the other, with the distribution of those that do neither the furthest left. The distributions of those that only do one or the other are relatively similar, with the distribution for firms

[17] In Figure 1.5 in Appendix 1.A we provide this graph defining importing intensively instead as being in the top 20th percentile within the same industry in terms of the ratio of imports from less developed countries to output, demonstrating that the same conclusions hold.

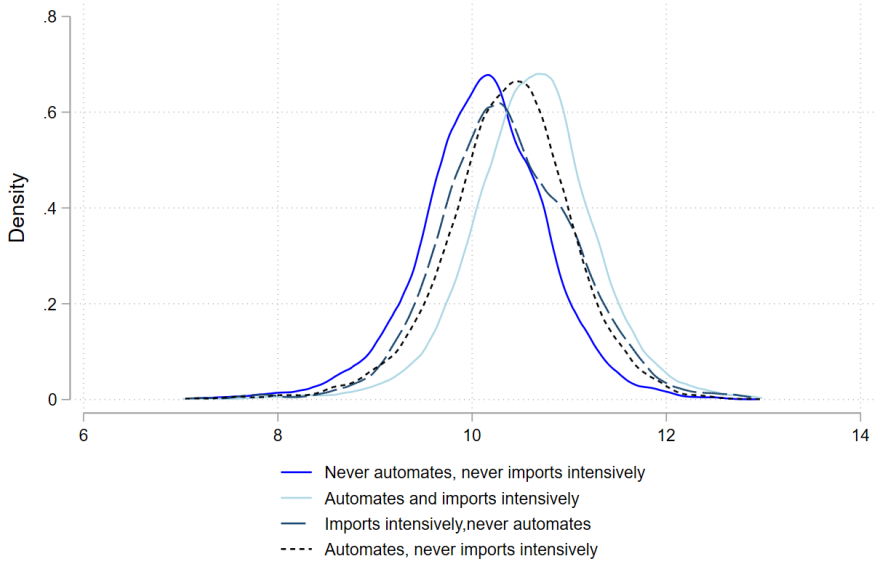


Figure 1.6: Labour productivity distribution by automation and importing from less developed countries

Notes: This figure displays the distribution of the log of value added per worker for the four groups of firms in terms of whether they automate using any of the three technologies at any point in the data and whether they import intensively from less developed countries, defined as being in the top 20% of the sample, at any point in the data.

that only automate being just a very slight shift right of that for firms that only import intensively. In Figure 1.10 in Appendix 1.A we provide the analogous figure for affiliates in less developed countries. Again the distribution of firms that both automate and have affiliates in less developed countries is the furthest right with the distribution for firms that do neither furthest left. However, with affiliates the distribution for firms that only have affiliates and never automate is further right than the distribution for firms that only automate.

Figure 1.11 in Appendix 1.A also displays the share of firms that do both, do either and do neither by industry. In terms of import intensity the industries with the highest share of firms doing both are 'Computer products, electronics and optical', 'Chemicals and pharmaceuticals', 'Other transport equipment', 'Basic metal products' and 'vehicles and accessories'. In terms of affiliates, three of those are also in the top five for having the highest share of firms do both, with 'Machinery and Equipment' and 'Beverage' also in the top 5. In terms of import intensity, the industry with the highest share of firms only importing intensively is 'Leather, fur and footwear', followed by 'Textiles and Clothing', as might be expected.

1.5 Empirical strategy

In Section 1.4 we demonstrated that firms that will select into automation in the future already have distinct characteristics relative to those that never automate and the gap is then wider for firms that have automated. This provides some early indication that automation might have a causal impact on firm outcomes, but could also be driven by differences between firms that automate at different times. In this section we aim to provide further evidence by attempting to evaluate the causal impact of automation on our key firm outcomes, focusing from this point onwards on robot use only. The baseline specification we aim to estimate is:

$$\mathrm{X}_{it} = \alpha + \beta \mathrm{Robots}_{it} + \theta \mathrm{industry\text{-}year}_{it} + \phi \mathrm{region\text{-}year}_{it} + \gamma_i + \epsilon_{it} \tag{1.23}$$

where X_{it} is the outcome of interest for firm i in year t, Robots_{it} is a dummy variable for firm use of robots at time t and γ_i, $\mathrm{industry\text{-}year}_{it}$ and $\mathrm{region\text{-}year}_{it}$ are firm,

industry-year and region-year fixed effects.[18] By including firm fixed effects our specification estimates the impact of changes in robot use within firms over time. The inclusion of industry-year fixed effects aims to soak up any trends that may have differentially affected industries in different time periods. We further include region-year fixed effects in order to control for any policy changes that affected specific regions in specific time periods.

1.5.1 Identification challenges

There are two major identification challenges when estimating the model in equation 1.23. First, this equation could suffer from reverse causality. In Section 1.4 we demonstrated that firms that select into automation in the future have distinct characteristics in terms of their international orientation, size and productivity, relative to firms that do not. It is possible that firm-level time-varying changes to importing or multinational activity could affect time-varying changes to robot adoption. There is recent evidence from Bernard et al. (2019), for example, that firms that offshore subsequently reorganise production, shifting activities at their headquarter firms towards more innovative and technology-focused production. If this also applied to the adoption of technologies, such as robots, and varied between firms within the same industry or region in a given year, then the demeaned robot use dummy variables could be correlated with the demeaned error term, resulting in biased estimates. Second, there could be unobserved firm-level time varying shocks that affect both robot adoption and the outcome variables. For example it is plausible that there could be firm-specific shocks to labour demand or trade that also affect the decision to automate.

In order to isolate the causal impact of robot use on firm performance and offshoring, we hence construct two separate instruments that should influence robot adoption over time but should not directly affect firm outcomes, such as importing or multinational activity. The first follows Webb (2019) and maps globally filed robotics patents to occupations conducting similar tasks. We then use the ex-ante

[18]The robot use dummy variables are provided every four years. In the case where the firm reports using or not using robots at both end of the four year period, we assume it continues to do the same in the years in between. In the case where the firm switches between the start and end of the period, we assume the switch occurs between the second and third intermediate years.

'exposure' of industry-region pairs to advances in robotics over time, based upon their employment composition in the 1980s. The second serves as a robustness check and follows Graetz and Michaels (2018). We outline these two methods below.

1.5.2 Exposure to advances in robotics using patents

One of the best ways of tracking technological progress in automation at its origin is through studying patents. We aim to develop a measure of the raw technological supply side for automation through looking at the stock of robot related patents, linked to the occupations that they potentially replace. Our motivation is that many technological developments stem from research conducted at universities and breakthroughs are often because of advances in complementary technologies in other fields or specific discoveries. For example, industrial robots have typically had limited mobility due to the cost of batteries and sensors. Only recently, cost declines for components such as lidar, which are used as sensors, or small and powerful batteries, have led to advances in mobile robotics. Cost declines for lidar were largely driven by the security industry and battery technology due to smartphones.

For manufacturing firms in Spain, these advances which made it suddenly possible to purchase a good quality robot at a low enough price were largely exogenous to them. Different manufacturing processes also involve different tasks, some of which became possible to automate with robots, while others did not. In the IFR data, most of the robot applications are conducted by a robotic arm, while smaller, intricate processes involving manual dexterity are still relatively hard to automate. These patterns are plausibly exogenous to the firm and particularly to a Spanish firm, given most robotics research and development generally occurs outside of Spain.

This method builds upon these observations and aims to measure the stock of robotics inventions aimed at conducting specific tasks. A detailed summary of the method can be found in Webb (2019), although we differ in this paper by adding a time dimension. The method uses the text of patents to identify the tasks that

robots can perform, then quantifies the extent to which each occupation in the economy involves performing similar tasks. A brief summary is as follows:

The first step of this method involved selecting patent publications related to industrial robots from the Google Patents Public Data database using keyword searches of patent titles and abstracts, and CPC codes. The next step involved extracting verb-noun pairs from the patent titles. This was done use a dependency parsing algorithm (Honnibal and Johnson, 2015) to determine the syntactic relations of the words in the sentence. For each verb, its direct object was identified by the algorithm, if it existed. The verb and noun were then lemmatized, so that, say, 'predicting' and 'predicted' were both recorded as 'predict'. Stop words such as 'use' and 'have', which do not express economic applications, were dropped.

The O*NET database was then used as the source of information on occupations. As noted above, each occupation consists of a collection of tasks described in natural language. The same dependency parsing algorithm was used to extract verb-noun pairs. Before calculating an exposure score for each verb-noun pair, nouns were grouped into conceptual categories. This was done using WordNet (Miller, 1995), a database developed at Princeton University that groups nouns into a hierarchy of concepts. For example, the ancestors of 'economist' are 'social scientist', 'scientist', 'person', 'causal agent', 'physical entity', and 'entity'.

An occupation's final exposure score was then calculated using the set of aggregated verb-noun pairs extracted from its task descriptions and this score expresses the intensity of the stock of patenting activity up until that date directed towards the tasks in that occupation. We hand construct a crosswalk between the O*NET SOC codes and ISCO 68 codes.[19] We then develop an industry-region level measure of 'exposure' using the 1980s occupational composition of employment in a given region and industry and the patent measures for each occupation code. The instrument for robot use by firm i in industry j and region r at time t is then:

$$\text{robot}_{irjt} = \text{exposure}_{rjt}^{robotpatent} \tag{1.24}$$

[19]We construct our own crosswalk because we could not find an established crosswalk from O*NET to ISCO 68 without combining multiple different crosswalks to go from O*NET to US SOC to ISCO88 then to ISCO68. We are happy to share this crosswalk and make it available on request.

where

$$\text{exposure}_{rjt}^{robotpatent} = \sum_{o=1}^{n} \text{automatable}_{ot}^{patent} \times \text{employment}_{orj1981} \tag{1.25}$$

where $\text{automatable}_{ot}^{patent}$ is the patent measure of automatability for occupation code o in year t and $\text{employment}_{orj1981}$ is the employment share of occupation o in region r and industry j in 1981.

We estimate 1.23 in levels with firm fixed effects. Because our instrument is at the industry-region-year level, there is insufficient variation across industries within different regions to include both industry-year and region-year fixed effects and so we aggregate the NACECLIO industries to a higher level industry grouping and include industry group-year fixed effects, along with region-year fixed effects.[20]

Insights from exposure measures

Most and least exposed occupations: Table 1.17 in Appendix 1.B displays the ISCO 68 occupations with the highest and lowest exposure scores scores. A more detailed discussion on the original O*NET SOC exposure scores can be found in Webb (2019). The most exposed occupations include various kinds of materials movers in factories and warehouses, and tenders of factory equipment. The most exposed category is 'Material-Handling and Related Equipment Operators, Dockers and Freight Handlers'. The least-exposed occupations include clergy, government administrators, accountants, jurists and authors. These do not primarily involve the kinds of repetitive manual tasks that robots automate.

Most and least exposed industries and regions: Figure 1.7 displays the most and least exposed industry-region pairs, using the 1980s employment weighting of the occupational exposure scores. The most exposed industry-region pair is the

[20]These groups are: 1. Heavy Manufacturing with the following industries included: industries Chemicals and pharmaceuticals, Plastic and rubber products, Nonmetal mineral products, Basic metal products, Fabricated metal products and Machinery and equipment and Timber. 2. Light manufacturing with the following industries included: Meat products, Food and Tobacco, Beverage, Textiles and Clothing, Leather, fur and footwear, Paper and Printing. 3. Complex manufacturing with te following industries included: Computer products, electronics and optical, Electric materials and accessories, Vehicles and accessories, Other transport equipment, furniture and Other manufacturing.

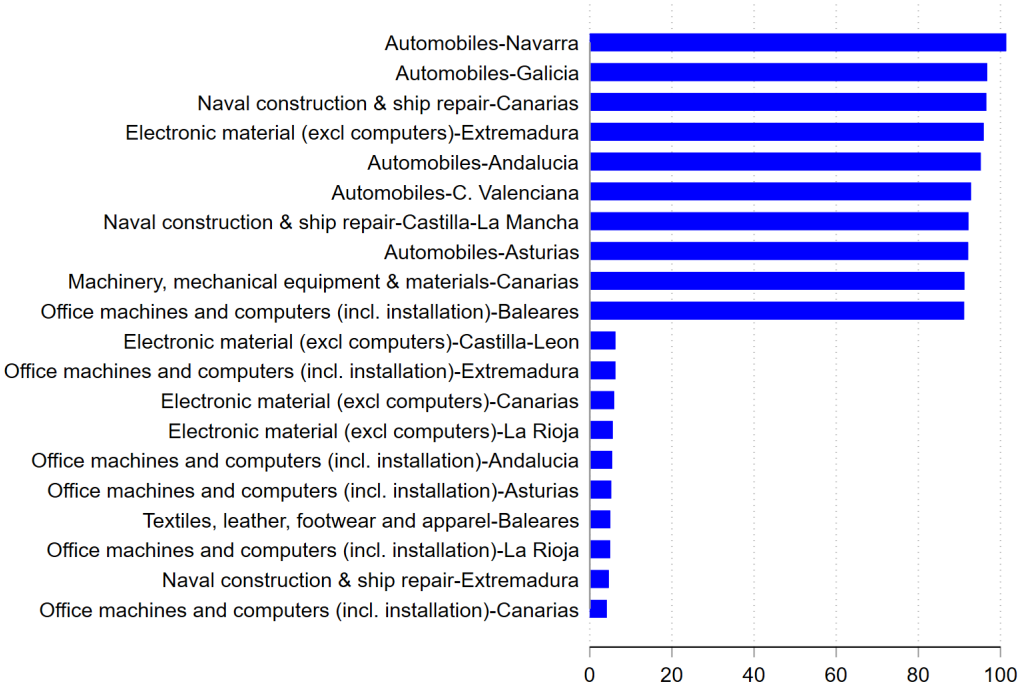


Figure 1.7: Top 10 industry-region pairs with highest and lowest median exposure

Notes: This figure displays the industry-region pairs with the highest and lowest median employment-weighted patent robot exposure scores across years.

construction of automobiles and spare parts in Navarra, while the least exposed industry-region pair is Office machines and Computers (including installation) in Canarias. We find that there is substantial variation in the exposure of the same industries in different regions, suggesting that the same industry can have a different occupational composition depending on its location.

1.5.3 Alternative exposure measure using robot sales data

For comparison we also construct an alternative measure for exposure to industrial robots using sales data from the IFR. The IFR provides global sales and operational stock of industrial robots by 'application'. Examples of the IFR applications are 'metal casting' or 'plastic moulding' or 'arc welding'. In a similar vain to Graetz and Michaels (2018) we hand match robot applications to ISCO 68 oc-

cupation codes.[21] We do so as follows: we define a 4 digit occupation code as 'automatable' in a given year using industrial robots if its title or formal description contains any of the words included in the application titles of the IFR data and there is a positive operational stock of that application robot type. We then develop an industry-region level measure of 'exposure' in a similar way to that used for patents by combining the 1980s occupational composition of employment in a region and industry and the classification of occupation codes. As this measure has limited time series variation, we additionally interact the exposure measure with the global operational stock of all industrial robots to obtain a shift share instrument.

1.5.4 Threats to identification

The identifying assumption for our patent-derived instrument is that the demeaned transformation of the 1981 occupational employment shares of industry-region pairs times time-varying occupation scores is uncorrelated with the demeaned error term. This would be the case if robotics inventions are indeed driven by the technology supply side and not demand side factors. One potential threat to identification would be if global robotics inventions over time are concentrated on tasks that are costly or difficult to offshore or unsafe for humans to perform and so there is more incentive to find ways to automate. An example of this could be tasks involving very large, heavy object that are expensive to move around, for example car chassis, or tasks involving dangerous procedures or handling chemicals, such as plastic moulding or metal casting. If this were the case then we might find that Spanish region-industry pairs with a high baseline employment share of occupations conducting these tasks, might, over time, experience both an increase in relevant global patenting and relatively little offshoring from Spain. This would lead to estimates of the impact of robots on offshoring that were biased downwards. Given we find a positive relationship, this should only serve to dampen our results.

[21] We construct these measures ourselves, rather than using the measures constructed by Graetz and Michaels (2018), because their measures are constructed for the US Census occupational classification and we would need to combine multiple crosswalks to convert these to ISCO 68, or construct a new crosswalk ourselves, which would result in a loss of accuracy relative to conducting the hand-matching ourselves.

Another possibility could be that international firms or research institutions have tried to invent robots specifically to automate the same occupations that Spanish firms have wanted to offshore. This could be the case if these are particularly labour intensive occupations or occupations suitable to both automation and offshoring. If this were the case then we might expect that region-industry pairs with a high baseline employment share of occupations conducting these tasks, might, over time, experience both an increase in relevant robotics patenting and offshoring. If these changes occurred at the same time and within industry groups, then we might overestimate the impact of robot use on offshoring. Anecdotally, it does not appear that this is the case; as described above, many robotics inventions have been linked to robotics arms and robots are still incapable of many labour intensive tasks that are offshored. However, it remains a possibility and so we try to address this challenge in a few ways. By controlling for region-year and industry-group year fixed effects, we remove all but within-region over time and within-industry-group over time trends. We also add additional controls for factors that have made it easier to offshore certain tasks over time, that could have co-occurred with patenting. We control for global supply shocks that reflect the increased opportunities to offshore production to countries, such as China and Eastern Europe, via our Hummels et al. (2014) style instrument and also for changes in intermediate import tariffs set at the EU level that may have made it cheaper to offshore.

A final concern is around import competition. We have not explicitly included the role of import competition in our analysis so far, although this time period was also one of increased competition for Spanish firms, particularly from China. One common way of controlling for import competition in the literature is to use Autor et al. (2013)) style local labour market shift share instruments. In our specification we include region-year fixed effects, which would fully absorb such a control. Import competition would be a threat to our identification if it affected firms within the same industry group and region and influenced both global robotics patenting and Spanish firms' offshoring or firm performance at the same time. One way in which we are able to address this possibility is through controlling for exogenously set tariffs and we do so with our same-industry import tariff control. This control

aims to soak up industry-time varying changes that may have influenced import competition through lowering trade costs.[22]

1.6 Results

1.6.1 Baseline results

Table 1.5 presents the baseline results using the patent instrument. We take the IHS transformation of the value of imports, the import intensity and the number of affiliates, due to the prevalence of zeros. We log all other variables except for the shares. We cluster standard errors throughout at the industry-region-year level in light of the sampling design of the ESEE data and our instruments being at this level.[23] In all regressions the first stage is strong, demonstrating that the patent exposure measure is a good predictor of whether a firm uses robots. Panel (a) displays results for the outcome variables relating to imports and multinational activity involving less developed countries. Starting to use robots has a substantial positive impact on imports from less developed countries. a positive impact on the intensity of imports from less developed countries and a positive impact on the number of affiliates in less developed countries. Robot use has no statistically significant effect on the share of imports from less developed countries, however, or on the share of affiliates in less developed countries.

These coefficients imply substantial impacts of robot adoption. The coefficient for imports from less developed countries implies that starting to use robots causes an approximately 13 fold increase in the value of imports. There are a few reasons why this value might be so high. This coefficient accounts for both the extensive margin and intensive margin of trade, which we explore in the next section. For firms switching from not importing at all from less developed countries to starting to do so (which we find to be the case), the percentage change will be large. In addition, our robot adoption variable takes only the valuss of zero or one and so a

[22] One final additional concern could be around the validity of shift-share instruments. Recent literature stresses the importance of controlling for the shares used to construct the shift-share variables (see, for example, Borusyak et al. (2019)). In our specification, however, we include firm fixed effects and so this is not a concern; if we added such a control it would drop our of the regression because the shares are time-invariant.

[23] We also tried alternative methods for clustering, such as at the firm level, finding this did not change the significance level of any of the coefficients in our baseline regressions.

switch could imply anything from starting to use one robot to switching production to a fully automated system. If it is the latter we are mainly capturing, then we would expect large results. We further explore the impact of the intensive margin of automation in Section 1.6.5. We also try replacing the IHS transformation with alternatives in the Section 1.7.6, demonstrating that this result still holds when we do so. The coefficient for import intensity implies that starting to use robots causes approximately a 5% increase in the value of imports relative to output, although this result is only weakly significant. The coefficient for affiliates implies that starting to use robots causes firms to increase their number of affiliates in less developed countries by approximately 50%.

Panel (b) displays the results for the domestic outcome variables. We find that, when using this IV, starting to use robots raises employment, the value of goods and services produced, labour productivity and TFP and decreases the labour share. The effect of robot use on domestic outcomes is also substantial. These results imply that within-firm switches to using robots increase employment in the firm in Spain by around 60%, cause a more than twofold increase in output, decrease the labour share of income by 0.49, more than double labour productivity and cause a more than threefold increase in TFP. Table 1.22 in Appendix 1.B also includes this and all other results relating to TFP for Olley-Pakes TFP instead.

We also repeat these baseline results with weightings by baseline firm output. Table 1.20 in Appendix 1.B displays these results. Weighting the observations does not qualitatively change our conclusions, except for employment, which no longer has a statistically significant coefficient. The remaining coefficients are all, in fact, larger in magnitude with weightings included. In addition, we repeat this analysis keeping only firms which remain in the sample for at least 5 years, finding that doing so barely affects the results, which we report in Table 1.19 in Appendix 1.B.

Table 1.18 in Appendix 1.B also presents the baseline results without instrumenting for robot use. These results suggest that without using an IV, within-firm switching into using robots is associated with an increase in the value of imports from less developed countries and the intensity of imports from less developed countries, higher employment, greater production of goods or services, higher labour productivity and higher TFP. There are two main differences between the

TABLE 1.5. BASELINE IV RESULTS

	IV FIXED EFFECTS REGRESSIONS 1990-2016				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports	Import intensity	Import share	Affiliates	Affiliate share
	(1)	(2)	(3)	(4)	(5)
Robot use	12.4***	0.049*	-0.14	0.49**	0.12
	(3.09)	(0.026)	(0.16)	(0.25)	(0.11)
First stage coef.	0.0051***	0.0051***	0.0049***	0.0041***	0.0041***
	(0.00077)	(0.00077)	(0.00076)	(0.00100)	(0.00100)
First stage F stat	43.5	43.5	40.5	17.1	17.1
Observations	37285	37256	37517	23067	23045
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment	Output	Labour share	Labour productivity	TFP
	(6)	(7)	(8)	(9)	(10)
Robot use	0.60**	2.49***	-0.49***	1.51***	3.06***
	(0.27)	(0.45)	(0.11)	(0.33)	(0.53)
First stage coef.	0.0048***	0.0049***	0.0048***	0.0047***	0.0053***
	(0.00079)	(0.00077)	(0.00080)	(0.00080)	(0.00085)
First stage F stat	36.5	40.2	35.8	34.9	39.1
Observations	35910	37371	35845	35523	32025
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-10. The variables in columns 6, 7, 9 and 10 are logged, while the variables in columns 1,2 and 4 are transformed by IHS. The instrument used is the patent derived measure of industry-region-year exposure to robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

OLS and IV results. First, the magnitude of the coefficients for OLS is far lower in all specifications. Second, for OLS there is no statistically significant relationship between robots and the number of developing country affiliates or the labour share. For the latter we provide a few guesses as to why this could be the case. It is plausible that firms that are looking to cut labour costs and so decreasing their labour share through means other than robots are then less likely to adopt robots, because they have already reduced labour costs through these means, such as outsourcing, hiring gig economy workers or investing in other forms of technology. Likewise, firms that have just opened affiliates abroad, particularly if for labour cost saving reasons, may be less interested in adopting robots at home. These possibilities would both bias coefficients in the OLS regressions towards zero. For the former, we also speculate that similar forces could be at play, biasing the other coefficients towards zero, such as firms that have already started offshoring to less developed countries being subsequently less likely to adopt robots.

1.6.2 Extensive margin and changes to employment and activities at affiliates

We next explore how robot use affects the extensive margin of importing from, or opening affiliates in, less developed countries. Our results above in Table 1.5 demonstrated that there was a very high coefficient for the impact of robot use on the value of imports, which could be driven by firms switching into importing from developing countries as well as the value of imports conditional on importing. Table 1.6 first displays results for the impact of robot use on a dummy variable for whether a firm imports from less developed countries, in Column (1) and whether it has affiliates in less developed countries in Column (2). For importing, the coefficient is statistically significant to the 5% level, showing that within-firm switching into using robots has a positive impact on whether a firm imports from less developed countries. Starting to use robots makes firms approximately 57% more likely to import from less developed countries. For opening affiliates, the coefficient is statistically significant to the 10% level, showing a weak relationship between using robots and opening affiliates in less developed countries.

TABLE 1.6. EXTENSIVE MARGIN AND AFFILIATE ACTIVITIES

	IV FIXED EFFECTS REGRESSIONS 1990-2016					
	Extensive margin		Activities at main affiliate			
Dep variable:	Starts importing	Opens affiliates	Affiliate employment	Adaptation & assembly	Marketing & distribution	Similar production
	(1)	(2)	(3)	(4)	(5)	(6)
Robot use	0.57**	0.30*	2.77**	0.12	0.35**	0.42**
	(0.24)	(0.16)	(1.09)	(0.11)	(0.16)	(0.18)
First stage F stat	40.5	17.1	16.9	17.1	17.1	17.1
Observations	37517	23067	22631	23044	23044	23044
Firm FE	Y	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-10. All variables are dummy variables except for the affiliate employment variable, which is transformed by IHS. The instrument used is the patent derived measure of industry-region-year level exposure to robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

Columns 3-6 display outcomes at the firm's main foreign affiliate. Starting to use robots also increases employment at the firm's main affiliate quite substantially. In terms of the primary activity conducted at the main affiliate, we find a positive impact of robot use on firms starting to report that 'the activity of the affiliated company consists only in commercialization or distribution' or that 'the main affiliated company manufactures similar products to those manufactured in Spain by its company', while there is no impact on reporting that 'the main affiliated company carries out adaptation and/or assembly activities of parts provided by the company in Spain'. These results suggest that the increase in multinational activity is concentrated in horizontal, rather than vertical FDI, where horizontal is associated with duplicating similar production stages across countries, while vertical reflects locating different production stages in different countries. These results provide some suggestion that it is plausible that as firms start to automate, they still expand abroad, but their focus shifts away from labour cost differentials and more towards market expansion.

1.6.3 Sequencing of automation and importing intensively

Our model predicted that in light of the fixed costs involved with automating or offshoring, the sequencing of these decisions would matter for the impact of au-

tomation. If a firm offshored first, we would expect that the displacement effect would occur for offshore labour, as well as domestic labour, dampening the impact on domestic labour. If a firm had not offshored first, we would not expect a displacement effect for current offshore labour. In Table 1.7 we display our baseline results firstly restricting the sample to only firms that started importing intensively from less developed countries prior to using robots and then using the full sample with interaction terms. As in section 1.4, we define importing intensively from less developed countries as being in the top 20% of the sample in terms of imports from less developed countries, equivalent to importing 1% or more of total output. Given that multinational activity variables only start from 2000, limiting the sample this way is not feasible for multinational activity and so we focus on importing variables only.

Panel (a) displays the results for the limited sample. Focusing only on firms that started importing intensively from less developed countries first yields quite different results to those for the full sample. For this group, starting to use robots has no effect on the value of imports from less developed countries, the import intensity or on the extensive margin. Its only effect is to reduce the share of imports from less developed countries by approximately 0.39. These results could suggest that firms that were already importing intensively, upon adopting robots, do not change their imports from less developed countries, but start importing from high-income countries, meaning the less developed country share of imports declines.

Panel (b) displays the results when including the full sample but using the interaction term between a dummy variable for being a firm that imported intensively from less developed countries before adopting robots and robot use. For the value of imports, the coefficient on the interaction term for firms that imported intensively first is negative and significant, suggesting that there is a statistically significant difference between the impact of robots on imports from less developed countries depending on the sequencing of importing and robot adoption. The same applies for the extensive margin. For the other two variables, however the difference is not statistically significant. These results generally support our model prediction that the sequencing of these decisions matters for their impact. In Figure 1.21 we also show that these results hold if we instead define importing

TABLE 1.7. SEQUENCING OF IMPORTING INTENSIVELY AND ROBOT USE

	IV FIXED EFFECTS REGRESSIONS 1990-2016			
	Panel (a): Only firms that import intensively first			
Dep variable:	Imports (1)	Import intensity (2)	Import share (3)	Starts importing (4)
Robot use	0.49	0.037	-0.39**	-0.26
	(3.17)	(0.051)	(0.20)	(0.25)
First stage F stat	29.5	29.3	27.3	27.3
Observations	5612	5602	5665	5665
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Region-year FE	N	N	N	N
Industry Group-Year FE	N	N	N	N
	Panel (B): Full sample with first importing intensively interaction			
Dep variable:	Imports (5)	Import intensity (6)	Import share (7)	Starts importing (8)
Robot use	16.0***	0.051	-0.18	0.76**
	(4.05)	(0.033)	(0.21)	(0.31)
Robot * first import intensively	-14.3***	-0.0085	0.16	-0.76***
	(3.62)	(0.028)	(0.18)	(0.27)
First stage F stat	18.5	18.4	17.7	17.7
Observations	37285	37256	37517	37517
Firm FE	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-8. Importing intensively is defined as being in the top 20% of the sample in terms of import intensity from less developed countries, defined as the value of imports from less developed countries scaled by output. Panel (a) limits the sample to only firms that import intensively from less developed countries before adopting robots. Panel (b) looks at the interaction between robot use and whether a firm first imported intensively from less developed countries before adopting robots, or not. The dependent variables of value of imports and import intensity are transformed by IHS. The instrument used is the patent derived measure of industry-region-year level exposure to robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

intensively as being in the top 20% within industries, or as the top 15% of the sample, although for the latter, for the limited sample regression, the significance of the result for the import share drops out.

1.6.4 Changes to employment in Spain

We next delve into the details of what was happening in terms of employment within the firms in Spain. Our results in Table 1.5 showed a strong positive impact of robot use on output, a strong negative impact on the labour share and a weak positive impact on employment in Spain. These results suggest that the impact on the labour share is not through decreased employment but through total income increasing relatively more than labour costs. We begin by picking apart this result for the labour share by additionally looking at how robot use affected total labour costs. Column 1 in Table 1.8 shows that the impact on labour costs was similar to the impact on employment.

Our model predicted that automation has both a productivity effect, whereby firms that automate can benefit from lower marginal costs, increase productivity and expand, and a displacement effect whereby a subset of tasks are automated. This would suggest that we should see the reorganisation of production within firms in Spain as some tasks are displaced and the composition of employment shifts to the tasks that are not feasible to automate. The ESEE data asks firms whether 'during the financial year, there has been a significant change in the number of regular workers', allowing us to evaluate if robot use does lead to this composition shift. Column 2 shows that starting to use robots does indeed have an impact on the changing of regular workers, with robot use leading firms to be around 44% more likely to report a change. Columns 3-6 explore the impact on different types of employment. The category that is most affected is the number of engineers or graduates, which are more than doubled as a consequence of adopting robots and production workers, which are increased by around 80%. Meanwhile we find no effect of robot use on the number of administrative workers or non-graduates.

TABLE 1.8. CHANGES TO EMPLOYMENT AND OPERATIONS IN SPAIN

	IV FIXED EFFECTS REGRESSIONS 1990-2016					
Dep variable:	Labour Costs	Change regular workers	Production	Clerical	Nongraduate	Engineers & graduates
	(1)	(2)	(3)	(4)	(5)	(6)
Robot use	0.62**	0.44**	0.80**	-0.14	0.33	1.13***
	(0.27)	(0.19)	(0.31)	(0.35)	(0.37)	(0.41)
First stage F stat	36.0	40.7	36.7	36.7	35.1	36.3
Observations	35938	33124	35898	35898	35699	35719
Firm FE	Y	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-6. The labour costs variable is logged, while the variables in columns 3-6 are transformed by IHS. The instrument used is the patent derived measure of industry-region-year level exposure to robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

1.6.5 Intensive margin of automation

So far our results have focused on the extensive margin of robot adoption using the dummy variable for whether firms use robots or not. It is also interesting to consider how the intensive margin of the extent of robot use affects our outcome variables. Our dataset does not include the total number of robots used, but does provide one variable of relevance: the share of annual investment that is in the 'purchase of technical facilities, machinery and tools'. Combined with annual total investment and the average depreciation rate, we use this variable to construct a measure of the annual stock of technical facilities, machinery and tools, using the perpetual inventory method. We then interact this variable with the robot dummy variable to obtain a proxy for the stock of industrial machinery for firms using robots.

Table 1.9 shows the results using this intensive margin robot stock variable as the independent variable. Although the coefficient magnitudes now have a different interpretation, the results are qualitatively similar. An increase in the stock of industrial machinery has a positive impact on the value of imports from less developed countries, the intensity of imports from less developed countries and the number of affiliates in less developed countries. In terms of magnitude, a 10% increase in the stock of industrial machinery for robot users increases imports from less developed countries by around 7%, the intensity of imports from less devel-

TABLE 1.9. INTENSIVE MARGIN OF AUTOMATION

	IV FIXED EFFECTS REGRESSIONS 1990-2016				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports	Import intensity	Import share	Affiliates	Affiliate share
	(1)	(2)	(3)	(4)	(5)
Log machinery stock	0.73***	0.0031**	-0.00024	0.034**	0.0082
	(0.17)	(0.0015)	(0.0093)	(0.017)	(0.0071)
First stage coef.	0.091***	0.091***	0.088***	0.064***	0.064***
	(0.013)	(0.013)	(0.013)	(0.016)	(0.016)
First stage F stat	52.5	52.7	49.5	16.0	16.0
Observations	35216	35195	35435	22119	22097
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment	Output	Labour share	Labour productivity	TFP
	(6)	(7)	(8)	(9)	(10)
Log machinery stock	0.027*	0.13***	-0.030***	0.089***	0.19***
	(0.016)	(0.024)	(0.0061)	(0.019)	(0.034)
First stage coef.	0.085***	0.088***	0.085***	0.082***	0.087***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.014)
First stage F stat	43.3	49.2	43.6	40.6	38.8
Observations	34185	35336	34146	33846	30716
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the interaction term for the dummy variable for firm use of robots $\times$ the log of the stock of investment in industrial machinery and equipment and the dependent variables, listed in the columns 1-10. The variables in columns 1,2 and 4 are transformed by IHS, while the variables in columns 6, 7, 9 and 10 are logged. The instrument used is the patent derived measure of industry-region-year level exposure to robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

oped countries by around 0.03% and the number of affiliates in less developed countries by around 0.3%. For domestic outcomes the results are also qualitatively similar for the intensive margin: the intensity of use of industrial machinery for robot users very weakly raises employment and strongly raises output, labour productivity and TFP, while decreasing the labour share.

1.7 Robustness

1.7.1 Controlling for changes to world export supply

Our instrument relies on both cross-sectional variation in industry-region level occupational employment shares in 1981 and on time series variation in occupation-specific patenting activity. One concern is that changes in patenting activity could be correlated with other time-varying shocks that may have affected importing or multinational activity with developing countries. In order to counteract this possibility, we include both industry group-year and region-year fixed effects. However, we are not able to include industry-year fixed effects because this reduces the variation in our instrument too much. There is therefore the possibility that there could be time-varying shocks to industries within the same industry group, within regions, and, if these were correlated with our outcome variables we would incorrectly attribute the effect to automation.

In Section 1.3.5 we outlined a control variable similar to the world export supply instrument for offshoring used in Hummels et al. (2014), designed to encompass a firm's exposure to changes to the global supply side of offshoring that are external to Spain, but could potentially influence the firm's decision to offshore, for example, the rise of China as a global export hub or the rise of Bangladesh as a destination for apparel offshoring. This control variable is a plausibly exogenous determinant of offshoring and so we include it to ensure that our instrument for robot use is not in some way reflecting these global supply shocks, as opposed to only the technology supply side of robotics inventions.

Table 1.10 repeats the regressions in the baseline results table, also including this control. We see that this variable does have a strong impact on offshoring; in Columns 1-3 it is statistically significant, suggesting that firms that were initially importing from destinations that subsequently experienced an increase in their

TABLE 1.10. CONTROLLING FOR CHANGES TO WORLD EXPORT SUPPLY

	IV FIXED EFFECTS REGRESSIONS 1990-2016				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports (1)	Import intensity (2)	Import share (3)	Affiliates (4)	Affiliate share (5)
Robot use	13.4***	0.053**	0.042	0.50*	0.11
	(3.12)	(0.026)	(0.16)	(0.26)	(0.11)
Offshoring instrument	1.12***	0.0052***	0.20***	0.0068	-0.0096
	(0.13)	(0.0011)	(0.0072)	(0.019)	(0.0089)
First stage F stat	44.2	44.2	41.2	15.8	15.8
Observations	37285	37256	37517	23067	23045
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment (6)	Output (7)	Labour share (8)	Labour productivity (9)	TFP (10)
Robot use	0.59**	2.54***	-0.51***	1.54***	3.15***
	(0.27)	(0.46)	(0.11)	(0.33)	(0.55)
Offshoring instrument	-0.014	0.054***	-0.022***	0.029*	0.098***
	(0.010)	(0.019)	(0.0053)	(0.015)	(0.033)
First stage F stat	37.1	40.8	36.3	35.4	38.3
Observations	35910	37371	35845	35523	32025
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-10. The variables in columns 6, 7, 9 and 10 are logged, while the variables in columns 1,2 and 4 are transformed by IHS. The instrument used is the patent derived measure of industry-region-year level exposure to robots. The offshoring instrument is the firm baseline import share & IO table weighted value of industry level exports from the world to the four EU countries with the closest GDP per capita to Spain, transformed by IHS. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

exports to other similar countries to Spain, saw an increase in their imports, import intensity and import share from less developed countries. This variable also has a positive impact on output and TFP and a negative impact on the labour share. Including this control has little effect on the coefficients for robot use, except to marginally increase the magnitude or significance of the coefficients for import value, import intensity, output, the labour share and TFP variables and marginally decrease the significance of the coefficient for the affiliates variable. These results generally suggest that omitting this control for external factors that made firms more likely to offshore has not biased our results for the impact of automation, except perhaps very slightly downwards.

TABLE 1.11. CONTROLLING FOR IMPORT TARIFF CHANGES

	IV FIXED EFFECTS REGRESSIONS 1990-2016				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports	Import intensity	Import share	Affiliates	Affiliate share
	(1)	(2)	(3)	(4)	(5)
Robot use	12.7***	0.050**	-0.17	0.50**	0.12
	(2.99)	(0.025)	(0.16)	(0.24)	(0.10)
Log imported input tariffs	-0.079	-0.00037	-0.024	-0.0015	-0.00061
	(0.29)	(0.0020)	(0.015)	(0.019)	(0.0092)
Log same industry tariffs	0.73***	0.0047***	0.037***	0.0070	-0.00038
	(0.22)	(0.0014)	(0.011)	(0.014)	(0.0060)
First stage F stat	46.9	46.9	43.8	18.4	18.4
Observations	37241	37212	37473	23023	23001
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment	Output	Labour share	Labour productivity	TFP
	(1)	(2)	(3)	(4)	(5)
Robot use	0.63**	2.45***	-0.49***	1.48***	3.02***
	(0.26)	(0.43)	(0.11)	(0.32)	(0.51)
Log imported input tariffs	0.021	-0.048	0.0069	-0.046	-0.060
	(0.021)	(0.045)	(0.011)	(0.032)	(0.058)
Log same-industry tariffs	0.0022	0.050	-0.0071	0.037	0.056
	(0.015)	(0.033)	(0.0075)	(0.024)	(0.040)
First stage F stat	39.7	43.4	39.0	37.9	42.0
Observations	35866	37327	35801	35479	31981
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-10. The variables in columns 6, 7, 9 and 10 are logged, while the variables in columns 1,2 and 4 are transformed by IHS. The instrument used is the patent derived measure of industry-region-year level exposure to robots. The import tariff controls are the log of the trade-weighted average intermediate import and import tariffs. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

1.7.2 Controlling for import tariff changes

In the context of offshoring, a further factor we might be particularly concerned about is the fall in import tariffs that may have reduced the cost of offshoring to less developed countries, for example China. In this section we therefore include additional controls for industry level import tariffs. We do so by following the method of Chen and Steinwender (2019). Import tariffs for Spain are set at the EU level and are hence plausibly exogenous to Spanish firms. We construct two control variables: one for MFN import tariffs on the intermediate inputs of an industry, which we denote as 'imported input tariffs', and the other for MFN import tariffs on products in the same industry, which we denote as 'same-industry im-

TABLE 1.12. EXCLUDING THE FINANCIAL CRISIS

	IV FIXED EFFECTS REGRESSIONS 1990-2016 EXCLUDING 2007-2010				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports (1)	Import intensity (2)	Import share (3)	Affiliates (4)	Affiliate share (5)
Robot use	12.0***	0.037	-0.061	0.51*	0.040
	(3.59)	(0.029)	(0.18)	(0.26)	(0.11)
First stage F stat	33.8	33.8	33.8	17.0	17.0
Observations	31378	31349	31378	16918	16896
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment (6)	Output (7)	Labour share (8)	Labour productivity (9)	TFP (10)
Robot use	0.60*	2.67***	-0.52***	1.76***	3.20***
	(0.32)	(0.54)	(0.12)	(0.39)	(0.60)
First stage F stat	30.6	33.5	29.9	29.3	34.7
Observations	29771	31232	29706	29425	25998
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-10. The variables in columns 1,2 and 4 are transformed by IHS, while the variables in columns 6,7, 9 and 10 are logged. The instrument used is the patent derived measure of industry-region-year level exposure to industrial robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

port tariffs'. The construction of these variables was outlined in the Data section.

The baseline results adding these additional controls are displayed in Table 1.11. Adding these controls does not change our coefficients significantly. We find that, somewhat surprisingly, the intermediate input tariffs do not have a statistically significant effect on any of the outcome variables, while import protecting tariffs only have an impact on the import variables.

1.7.3 Excluding the Financial Crisis

The Global Financial Crisis fell within the time period of our analysis and so an additional concern is that this may have affected our results. We hence also repeat our analysis excluding the years from 2007-10. Table 1.12 repeats the analysis in Table 1.5 for the time period of only 1990-2006 & 2011-2016. The results change slightly but the only qualitative difference is that the coefficient on the import intensity variable is now not statistically significant.

TABLE 1.13. IFR INSTRUMENT FOR ROBOT ADOPTION

	IV FIXED EFFECTS REGRESSIONS 1990-2016				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports	Import intensity	Import share	Affiliates	Affiliate share
	(1)	(2)	(3)	(4)	(5)
Robot use	5.69**	0.048**	0.15	0.42*	0.028
	(2.23)	(0.020)	(0.17)	(0.23)	(0.10)
First stage coef.	0.021***	0.021***	0.021***	0.021***	0.021***
	(0.0030)	(0.0030)	(0.0030)	(0.0030)	(0.0030)
First stage F stat	51.8	51.2	51.8	16.2	16.2
Observations	40557	40530	40557	28358	28336
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment	Output	Labour share	Labour productivity	TFP
	(6)	(7)	(8)	(9)	(10)
Robot use	0.10	1.58***	-0.067	1.18***	1.73***
	(0.21)	(0.33)	(0.22)	(0.28)	(0.36)
First stage coef.	0.021***	0.021***	0.021***	0.021***	0.021***
	(0.0030)	(0.0030)	(0.0030)	(0.0030)	(0.0030)
First stage F stat	51.5	49.7	49.7	47.2	48.0
Observations	40454	40422	40419	39992	39184
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-10. The variables in columns 1,2 and 4 are transformed by IHS, while the variables in columns 6,7, 9 and 10 are logged. The instrument used is the IFR derived measure of industry-region-year level exposure to industrial robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

1.7.4 Alternative instrument using IFR data

Table 1.13 displays the results for the full time period repeated using the IFR exposure measure instead of the patent measure. The first stage coefficient and F stat demonstrate that this instrument is also a strong predictor of robot use. The results are qualitatively similar, with the main difference being that the coefficients on domestic employment and the labour share are now not statistically significant, suggesting that these results are less robust than for the other outcome variables. The coefficient for the value of imports from less developed countries is now slightly less statistically significant at only the 5% level and with a lower magnitude. The coefficients on output, labour productivity and TFP remain statistically significant at 1% but now have lower coefficients. We can also use this instrument to conduct a GMM difference test (a.k.a. C test) for the orthogonality of our two instruments. We are able to reject the null for both instruments, supporting their validity. Taken together these results generally suggest that the key results for the impact of robot use on imports from less developed countries, output, labour productivity and TFP are all robust to using an alternative IV for robot use.

1.7.5 Is this just the automotive sector?

The industry in our dataset with the highest prevalence of robot use is Vehicles and Accessories, with 62% of firm-year observations reporting that they use robots, relative to the mean for the full sample of 27%. An additional concern is that our results are driven solely by this one outlier industry, that is particularly reliant on robots for production. We hence also try removing this sector from our analysis. Table 1.14 displays the results. We find that when we exclude this sector, the impact of robot use on developing country imports remains statistically significant and the coefficient large in magnitude, but the significance and magnitude are somewhat lower than when autos are included. For domestic outcomes, on the other hand, the coefficients are all larger in magnitude and with a qualitatively different result for employment, which is now strongly significant and with a magnitude over twice as large. These results generally suggest that for industries other than autos, the impact of robots on domestic outcomes is greater, but the impact on

TABLE 1.14. EXCLUDING THE AUTOMOTIVE SECTOR

	IV FIXED EFFECTS REGRESSIONS 1990-2016				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports (1)	Import intensity (2)	Import share (3)	Affiliates (4)	Affiliate share (5)
Robot use	8.76**	0.062*	-0.095	0.59*	0.062
	(3.47)	(0.033)	(0.22)	(0.31)	(0.12)
First stage F stat	26.9	26.8	24.7	13.9	13.9
Observations	35432	35403	35653	21858	21836
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment (6)	Output (7)	Labour share (8)	Labour productivity (9)	TFP (10)
Robot use	1.55***	4.22***	-0.70***	1.90***	4.20***
	(0.37)	(0.85)	(0.18)	(0.49)	(0.90)
First stage F stat	22.9	24.3	22.3	23.0	23.7
Observations	34117	35512	34053	33748	30443
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-10. The variables in columns 1,2 and 4 are transformed by IHS, while the variables in columns 6,7, 9 and 10 are transformed by the Inverse-Hyperbolic Sine (IHS) transformation. The instrument used is the patent derived measure of industry-region-year level exposure to industrial robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

international outcomes weaker, perhaps reflecting the relatively high international orientation of the auto industry.

1.7.6 Alternative specifications for value of imports

We noted in our baseline results table that the coefficient for the value of imports from less developed countries was very high and the estimation problem of having many firm-year observations with zero imports from less developed countries. We hence also try repeating the analysis for this outcome variable replacing the IHS transformation with log(1+ x) and a PPML model. For the PPML model, given we have an IV specification we use the control function approach to estimate the second stage and then bootstrap the standard errors to adjust for having included a derived regressor. Table 1.15 displays the results. We find that trying these two alternative specifications does not change the magnitude of the coefficient substantially. The specification with the lowest magnitude coefficient is the PPML model, suggesting that there are some signs that selection into importing from less developed countries is inflating our coefficients, although the coefficient still remains

TABLE 1.15. ALTERNATIVE SPECIFICATIONS FOR IMPORTS REGRESSIONS

IV FIXED EFFECTS REGRESSIONS 1990-2016				
Dep variable:	Imports from less developed countries			
Specification:	IHS (1)	ln(1+x) (2)	PPML (3)	ln(x) (4)
Robot use	12.4*** (3.09)	11.9*** (2.95)	11.6** (4.86)	0.62 (1.36)
First stage F stat	43.5	43.5	-	10.6
Observations	37285	37285	37285	8464
Firm FE	Y	Y	Y	Y
Year FE	N	N	N	Y
Region-year FE	Y	Y	Y	N
Industry Group-Year FE	Y	Y	Y	N

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the value of imports from less developed countries, with the specifications in columns 1-4. The instrument used is the patent derived measure of industry-region-year level exposure to industrial robots. These results are not weighted. Standard errors for columns 1,2 and 4 are robust to heteroskedasticity and serial correlation at the industry-region-year level. For column 3 we use the control-function method to estimate the IV-PPML model and bootstrap the standard errors.

very high in magnitude at 11.6.

We also include the results taking only the log. All observations with zero imports from less developed countries then drop out of the sample, resulting in a far smaller sample size and a weak first stage. The result is not significant for this logged specification, further suggesting that when we focus only on changes to importing amongst firms already importing from less developed countries, there is less of an impact of robots, perhaps because an important part of the effect is the extensive margin of switching into importing.

1.8 Conclusion

In this paper we take advantage of a rich dataset of Spanish manufacturing firms between 1990 and 2016 to shed new light on the consequences of automation in a high-income country for imports and multinational activity involving less developed countries. In order to evaluate the causal impact of automation on trade and multinational activity, we exploit supply-side advances in the capabilities of robots over time that made it technically feasible to automate specific tasks. We follow Webb (2019) and use the text from patent titles to identify the tasks that robot-related patents substitute for, mapped to the occupations frequently conducting those tasks. This allows us to construct time-varying measures of exposure to au-

tomation with robots.

Using this instrument, we show that starting to use robots in Spain caused a within-firm increase in the value of imports from, and the number of affiliates in, less developed countries. Robot use also had a positive impact on the extensive margin of trade and multinational activity, leading firms to start importing from, or start opening affiliates in, less developed countries. For multinational activity, the expansion appears to be more directed at horizontal, rather than vertical FDI, however, with robot use increasing the probability that the primary activity of the main affiliate is marketing or distribution of products manufactured in Spain, or producing similar products to those manufactured in Spain, as opposed to assembly or adaptation of inputs supplied by the firm in Spain. We show that these results hold after controlling for global supply shocks that made it easier to offshore production, import tariff changes, excluding the period of the Global Financial Crisis and excluding the automotive sector. For comparison, we also construct an additional instrument for use of robots using the capabilities of industrial robots as identified from robot sales data from the IFR, in a similar way to Graetz and Michaels (2018). We find similar results using this instrument.

We further demonstrate that the sequencing of automation and offshoring matters for the impact of automation. When we focus on firms that were importing intensively from less developed countries before they starting to use robots, we find that robot use had no effect on the value of imports but decreased the share of imports sourced from developing countries, suggesting that automation does, to some degree, shift economic activity away from developing countries, but only for firms that had offshored production first. By contrast, for firms that automate first, starting to use robots has a positive impact on the probability that they import from less developed countries and the value of their imports.

Taken together, these results suggest that fears of the consequences of automation for 'reshoring' seem to have been overblown and have over-simplified what is, in fact, a complex relationship. In our sample of firms in Spain, more firms started to automate before they started to import intensively from, or open affiliates in, less developed countries. For these firms, automation generated new demand for imports and affiliates. In this paper we haven't been able to say anything about

the specific products imported or the countries affected. The increase in imports could have been in less labour-intensive goods, or concentrated in middle or upper middle income countries, such as China, rather than new offshoring destinations. The increase could also reflect imports from foreign firms that have also started automating or shifted to higher tech production processes. These are all important avenues for future research.

While we do find some negative consequences for the share of imports from less developed countries, particularly after the Financial Crisis, through the channel of firms that imported intensively before automating reducing their less developed country import share, we do not find any aggregate negative effects. Firms that imported intensively before automating did not decrease their absolute value of imports from less developed countries. For the full sample the total value of imports and the import intensity from less developed countries were also still increasing after the Financial Crisis. However, our modelling framework could imply that new entrant firms, that face an expanded task subset that is automatable, might exhibit different patterns. A further important avenue for future research is to consider whether this relationship has changed in more recent years or for new firms.

Acknowledgements Katherine Stapleton gratefully acknowledges funding from the ESRC and Oxford's CSAE to purchase the data for this paper. We would like to thank Abi Adams, Erhan Artuc, David Autor, Jan Bakker, Nick Bloom, Maria Bolgova, Doug Gollin, Clement Imbert, Beata Javorcik, Peter Neary, Matthias Qian, Pascual Restrepo, Tony Venables and Chris Woodruff, as well as participants of the Oxford Firms, Labour & Organisations Research Group, the Oxford Trade Workshop and the CSAE Workshop for insightful comments and feedback.

Appendices of Chapter 1

1.A Additional figures

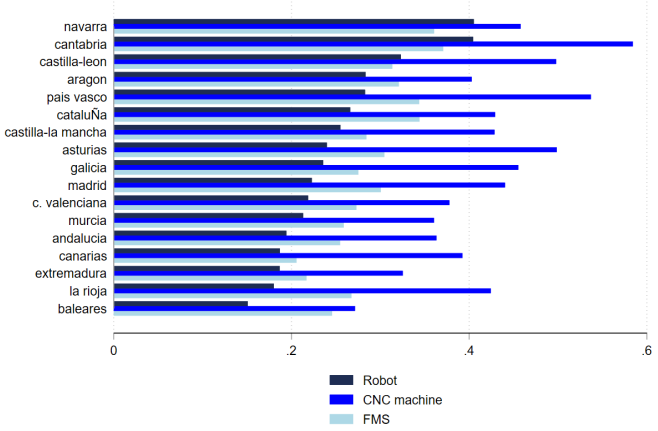


Figure 1.8: Technology use by region

Notes: This figure displays the share of firms using the three technologies, over all time periods, by Spanish region included in the ESEE data.

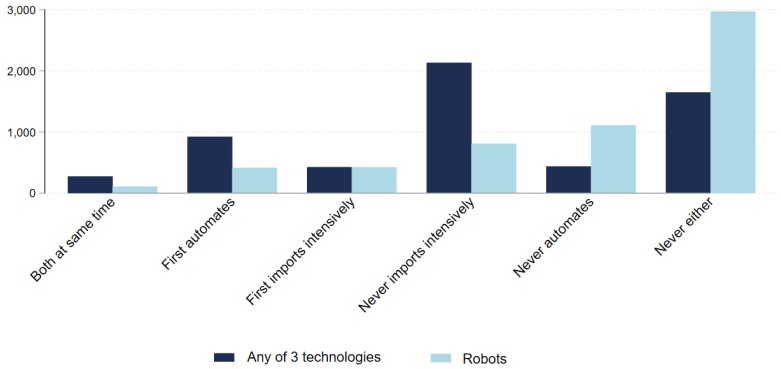


Figure 1.9: Sequencing of automation and importing intensively from less developed countries

Notes: This figure displays the number of firms by the sequencing of automation and importing intensively from less developed countries. Importing intensively is defined as being in the top 20% of the industry in terms of import intensity from less developed countries.

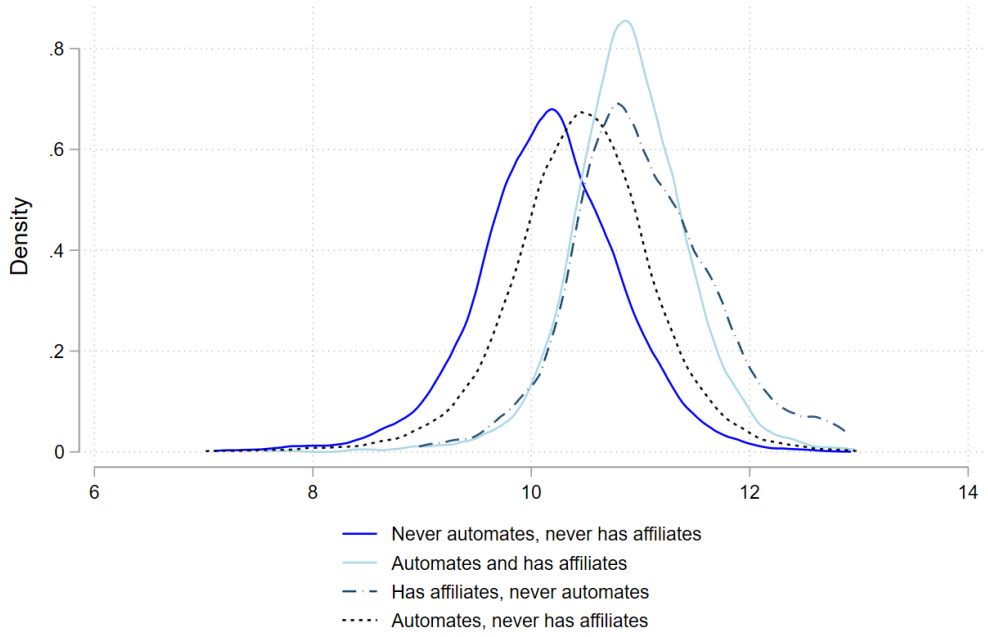
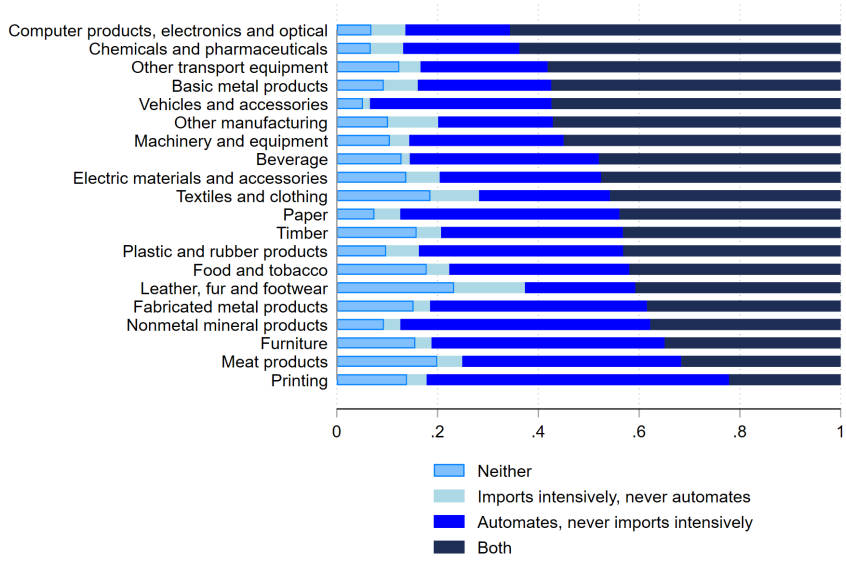
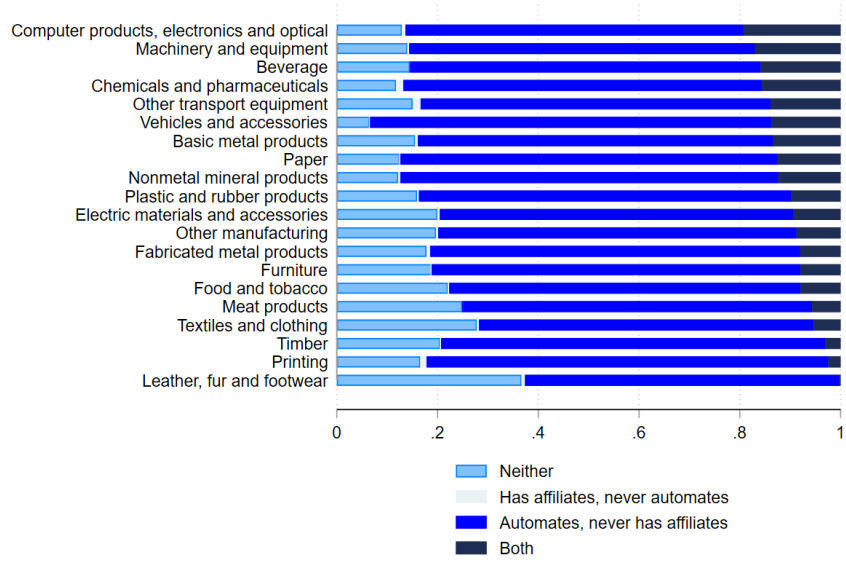


Figure 1.10: Labour productivity distribution by automation and less developed affiliate status

Notes: This figure displays the distribution of the log of value added per worker across firms depending on whether they ever have affiliates in less developed countries or ever use robots, neither or both.



(a) Importing intensively from less developed countries



(b) Affiliates in less developed countries

Figure 1.11: Automation, importing intensively or affiliates in less developed countries, by industry

Notes: Panel (a) displays the share of firms by whether they ever import intensively from less developed countries, defined as being in the top 20% of the sample or whether they ever automate, neither or both, by industry. Panel (b) displays the same in terms of whether they ever have affiliates in less developed countries.

1.B Additional tables

TABLE 1.16. CORRELATION BETWEEN TECHNOLOGY USE

	Robots	FMS	CNC Machines
Robots	1		
FMS	0.2261	1	
CNC Machines	0.3407	0.177	1

Notes: This table displays the correlation coefficients between the three dummy variables for use of robots, FMS and CNC Machines.

TABLE 1.17. ISCO 68 CODES WITH HIGHEST & LOWEST PATENT SCORES

Most exposed	Least exposed
97: Material-Handling & Related Equipment Operators, Dockers & Freight Handlers	14: Workers in Religion
98: Transport Equipment Operators	20: Legislative Officials & Government Administrators
84: Machinery Fitters, Machine Assemblers & Precision Instrument Makers	11: Accountants
54: Maids and Related Housekeeping Service Workers NEC	12: Jurists
96: Stationary Engine and Related Equipment Operators	15: Authors, Journalists & Related Writers

Notes: This table displays the top 5 and bottom 5 most and least robot exposed occupations at ISCO 68 level using the patent-derived exposure measure.

TABLE 1.18. RESULTS WITHOUT IV

	FIXED EFFECTS REGRESSIONS 1990-2016				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports	Import intensity	Import share	Affiliates	Affiliate share
	(1)	(2)	(3)	(4)	(5)
Robot use	0.38***	0.0021***	-0.0036	-0.0026	0.00091
	(0.083)	(0.00070)	(0.0052)	(0.0061)	(0.0029)
Observations	46435	46399	46435	28358	28336
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment	Output	Labour share	Labour productivity	TFP
	(6)	(7)	(8)	(9)	(10)
Robot use	0.053***	0.096***	-0.0058	0.031***	0.049***
	(0.0060)	(0.0077)	(0.0036)	(0.0086)	(0.0096)
Observations	44277	46220	44196	43688	39184
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in columns 1-10. The variables in columns 1,2, and 4 are transformed by IHS, while the variables in columns 6,7, 9 and 10 are logged. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

TABLE 1.19. BASELINE RESULTS EXCLUDING FIRMS WITH <5 YEAR OBS

	IV FIXED EFFECTS REGRESSIONS 1990-2016				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports	Import intensity	Import share	Affiliates	Affiliate share
	(6)	(7)	(8)	(9)	(10)
Robot use	12.4***	0.049*	-0.14	0.49**	0.12
	(3.11)	(0.026)	(0.17)	(0.25)	(0.11)
First stage F stat	43.1	43.1	40.1	17.1	17.1
Observations	37249	37220	37477	23067	23045
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment	Output	Labour share	Labour productivity	TFP
	(1)	(2)	(3)	(4)	(5)
Robot use	0.60**	2.50***	-0.49***	1.51***	3.06***
	(0.27)	(0.46)	(0.11)	(0.33)	(0.53)
First stage F stat	36.5	39.7	35.8	34.9	39.1
Observations	35896	37333	35831	35509	32025
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-10. Firms with less than five time series observations are excluded. The variables in columns 1,2 and 4 are transformed by IHS, while the variables in columns 6,7, 9 and 10 are logged. The instrument used is the patent derived measure of industry-region-year level exposure to industrial robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

TABLE 1.20. BASELINE IV RESULTS: WEIGHTED

	IV FIXED EFFECTS REGRESSIONS 1990-2016				
	Panel (a): International outcomes involving less developed countries				
Dep variable:	Imports	Import intensity	Import share	Affiliates	Affiliate share
	(1)	(2)	(3)	(4)	(5)
Robot use	16.5***	0.062**	-0.15	0.69*	0.19
	(4.25)	(0.031)	(0.18)	(0.35)	(0.14)
First stage F stat	30.7	30.8	30.7	11.5	11.5
Observations	37504	37481	37504	23064	23042
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y
	Panel (b): Domestic outcomes				
Dep variable:	Employment	Output	Labour share	Labour productivity	TFP
	(6)	(7)	(8)	(9)	(10)
Robot use	0.41	2.53***	-0.55***	1.72***	3.32***
	(0.31)	(0.52)	(0.13)	(0.41)	(0.64)
First stage F stat	27.2	30.9	27.2	26.4	30.5
Observations	35898	37371	35845	35523	32025
Firm FE	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-10. The variables in columns 6, 7, 9 and 10 are logged, while the variables in columns 1,2 and 4 are transformed by IHS. The instrument used is the patent derived measure of industry-region-year level exposure to robots. These results are weighted by baseline firm output. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

TABLE 1.21. CHANGING THRESHOLDS FOR IMPORT INTENSITY

	IV FIXED EFFECTS REGRESSIONS 1990-2016							
	Top 20% within industries				Top 15% of sample			
	Panel (a): Only firms that import intensively first							
Dep variable:	Imports	Import intensity	Import share	Starts importing	Imports	Import intensity	Import share	Starts importing
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Robot use	-0.65	0.051	-0.50**	-0.53*	0.58	0.036	-0.0086	-0.0027
	(3.48)	(0.055)	(0.22)	(0.30)	(4.54)	(0.074)	(0.26)	(0.34)
First stage F stat	22.4	22.2	20.2	20.2	14.1	14.0	13.1	13.1
Observations	5218	5209	5272	5272	4268	4259	4313	4313
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Region-year FE	N	N	N	N	N	N	N	N
Industry Group-Year FE	N	N	N	N	N	N	N	N
	Panel (B): Full sample with first importing intensively interaction							
Dep variable:	Imports	Import intensity	Import share	Starts importing	Imports	Import intensity	Import share	Starts importing
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Robot use	20.1***	-0.015	-0.33	1.14***	15.8***	-0.0069	-0.29	0.87***
	(5.90)	(0.044)	(0.26)	(0.43)	(4.36)	(0.036)	(0.21)	(0.33)
Robot * first imp intensively	-18.0***	0.046	0.22	-1.11***	-13.2***	0.046	0.24	-0.78***
	(5.14)	(0.037)	(0.22)	(0.37)	(3.63)	(0.028)	(0.17)	(0.27)
First stage F stat	12.2	12.1	12.0	12.0	17.6	17.5	17.2	17.2
Observations	37285	37256	37517	37517	37285	37256	37517	37517
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Region-year FE	Y	Y	Y	Y	Y	Y	Y	Y
Industry Group-Year FE	Y	Y	Y	Y	Y	Y	Y	Y

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the dependent variables, listed in the columns 1-8. columns 1-4 define importing intensively as being in the top 20% within the same industry in terms of import intensity from less developed countries. Columns 5-8 define it as being in the top 15% of the sample. The value of imports and import intensity are transformed by IHS. The instrument used is the patent derived measure of industry-region-year level exposure to robots. These results are not weighted. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

TABLE 1.22. OLLEY-PAKES TFP

	IV FIXED EFFECTS REGRESSIONS 1990-2016							
Instrument:	Patent							IFR
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Robot use	2.65***	2.97***	2.72***	2.60***	2.78***	3.34***		1.40***
	(0.48)	(0.60)	(0.50)	(0.46)	(0.54)	(0.76)		(0.32)
Log machinery stock							0.17***	
							(0.031)	
First stage F stat	39.1	30.5	38.3	42.0	34.7	23.7	38.8	48.0
Observations	32025	32025	32025	31981	25998	30443	30716	39184
Weighting	N	Y	N	N	N	N	N	N
World export supply control	N	N	Y	N	N	N	N	N
Tariff control	N	N	N	Y	N	N	N	N
Excluding 2007-2010	N	N	N	N	Y	N	N	N
Excluding automotive	N	N	N	N	N	Y	N	N

Notes: This table presents estimates of the relationship between the dummy variable for firm use of robots and the log of TFP estimated using the Olley-Pakes method, instead of Levinson-Petrin, for all of the main specifications and robustness checks. All regressions also include firm, region-year and industry group-year fixed effects. Standard errors are robust to heteroskedasticity and serial correlation at the industry-region-year level.

1.C Additional definitions

EU and OECD membership over time

The following countries were members of either the OECD or the EU in 1990: Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Japan, Luxembourg, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, Turkey, UK, US. These countries collectively account for 64% of Spain's total imports over the period of 1990-2016, according to UNCOMTRADE data.

The following countries switch status to join the group of OECD & EU countries during our sample period. Table 1.23 displays the countries that switched into this group during the sample, recorded by the year when they first switched into this group and the organisation they joined. These countries collectively account for 6% of Spain's imports over the period of 1990-2016. The countries in this group accounting for the largest trade volume are Mexico, with 1.3% of all imports from 1990-2016 and South Korea, with 0.9% of all imports.

TABLE 1.23. COUNTRIES THAT JOIN OECD & EU DURING SAMPLE PERIOD

Year	OECD	EU
1994	Mexico	
1995	Czech Republic	
1996	Hungary, Korea, Poland	
2000	Slovak Republic	
2004		Cyprus, Estonia, Latvia, Lithuania, Malta, Slovenia
2007		Bulgaria, Romania
2010	Chile, Israel	
2013	Croatia	

Notes: This table displays the countries that switched from the 'less developed country' group into the OECD & EU group during our sample, recorded by the date when they first joined either of these groups and the group they first joined. .

CHAPTER 2

Unravelling Deep Integration: Local Labour Market Effects of the Brexit Vote

with **Beata Javorcik, Ben Kett and Layla O'Kane**

Abstract This paper uses high frequency data on the near universe of job adverts posted online in the UK to study the impact of the Brexit referendum on labour demand between January 2015 and December 2019. We develop measures of local labour market exposure to the threat of trade barriers on both goods and services exports if the UK were to leave the EU without a trade deal. We find that regions that were more exposed to potential barriers on professional services exports saw a differential decline in online job adverts in the period after the referendum, particularly for higher-skilled jobs. This effect was distinct from the impact of the exchange rate depreciation, uncertainty surrounding future immigration policy and the threat of future barriers on trade in goods.

2.1 Introduction

The consequences of trade barriers for economic outcomes have been studied for a long time, yet the impact of uncertainty about possible future barriers, as opposed to barriers themselves, is less well understood. In addition, research that does exist on the impact of uncertainty about possible future barriers has typically focused on goods trade, rather than services trade, despite the growing role of services in global exports. With some of the world's major economies in a period of heightened uncertainty in their trade relations and the growing prevalence of deep trade agreements, which cover not just trade but other complex provisions aimed at deepening economic integration, this area of research is becoming increasingly important. The Brexit vote and its aftermath provide a unique opportunity to study the impact of trade policy uncertainty, particularly with regard to services trade.

This paper examines how the threat introduced by the Brexit vote of 'unravelling' decades worth of deep integration with the EU affected labour demand in the UK.[1] We use a high frequency dataset consisting of the near universe of job adverts posted online in the UK between January 2015 and December 2019, which allows us to evaluate the immediate labour demand response to the trade policy uncertainty caused by the referendum result and political events during the negotiation period. Figure 2.1 displays the normalised, smoothed time series of monthly job postings over the period of our analysis for both the UK and US. Both countries followed a similar trend before the referendum but started to diverge substantially from January 2018 onwards. In 2019 total job postings in the UK were 10% lower than in 2015, whereas in the US they were 64% higher.

We develop measures of local labour market exposure to potential future barriers on exports of goods and services to the EU. We take the 'threat' of no trade deal being reached between the UK and EU and evaluate the barriers that exporters would face under this scenario. For services, if the UK were to leave the

[1] Deep integration typically refers to economic integration that goes beyond trade to include freer movement of goods, services, capital, people and ideas. It is typically achieved through a Deep Trade Agreement (DTA) where DTAs are preferential trade agreements that cover not just trade but additional policy areas, such as international flows of investment or protection of intellectual property rights (Mattoo et al., 2020).

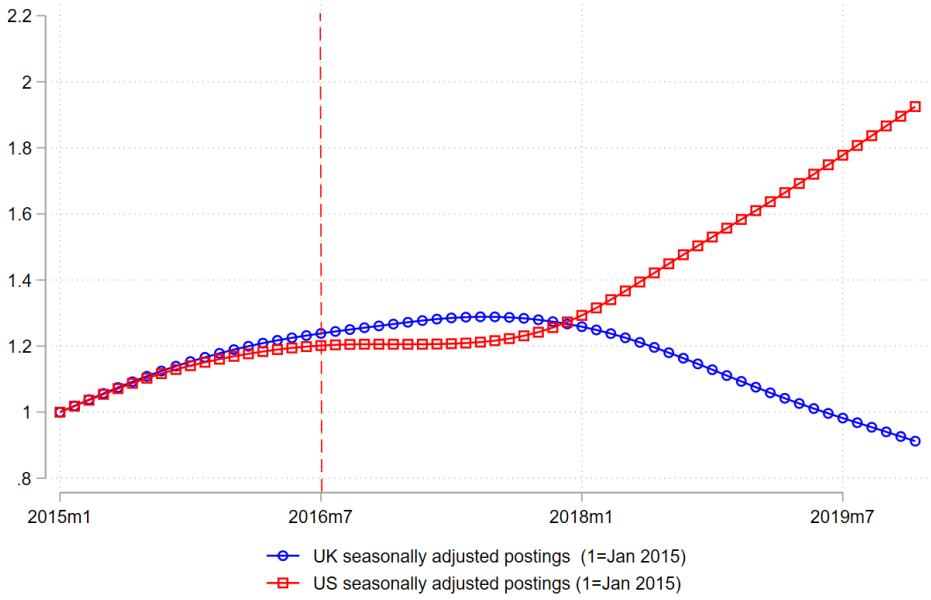


Figure 2.1: Monthly online job postings, UK and US

Notes: Both trend lines are smoothed using the Hodrick-Prescott time-series filter, removing cyclical components, and normalised to 1 in January 2015 for comparison. The vertical red dotted line identifies the date of the Brexit referendum.

EU and Single Market without a deal, it would become a third party to other EU & European Economic Area (EEA) countries and trade in services would revert to the terms set under the General Agreement on Trade in Services (GATS). The UK would then face rules set by individual member states.[2] This would represent a major shift in policy for many UK service sectors as the EU Single Market is a highly integrated area for trade in services, and in 2018 services accounted for 81% of UK economic output and 46% of exports (ONS, 2019).

We exploit the pre-vote variation in the difference between the regulations that individual member states place on services trade with other EEA members and non-EEA members, in different sectors. For each export destination and sector, we construct the regulatory 'gap' between the threat of the stringent restrictions UK exporters could face and the minimal restrictions they currently face using the OECD's Services Trade Restrictiveness Index (STRI) and Intra-EEA STRI, respec-

[2]The EU's Single Market is an area that seeks to guarantee the free movement of goods, capital, services, and labour within the EU. It encompasses the EU's 27 member states plus Iceland, Liechtenstein and Norway through the Agreement on the EEA and Switzerland through bilateral treaties.

tively, focusing primarily on trade barriers to professional services exports. UK industries were differentially exposed to this threat of regulatory barriers depending on the pre-vote composition of their exports by country. And, in turn, local labour markets were differentially exposed to this threat based upon their pre-vote employment composition by industry.

For goods, no trade deal with the EU would mean reverting from tariff-free trade to trade under World Trade Organisation (WTO) terms, where the UK would face Most Favoured Nation (MFN) tariffs on its exports to the EU. We exploit the product-level variation in EU MFN tariffs to construct measures of the tariff threat to sectors and local labour markets.[3]

We relate the monthly job advert data for each of the UK's 213 local labour markets, named 'Travel to Work Areas' (TTWAs), to the pre-vote exposure of these TTWAs to potential future barriers on exports to the EU, controlling for the impact of the exchange rate depreciation and exposure to immigration policy changes. We follow a difference-in-differences approach comparing job adverts in highly exposed localities to less exposed localities in the post-referendum relative to the pre-referendum period. We demonstrate that this strategy is warranted through several robustness checks, which show no evidence of differences in pretrends. The result of the Brexit referendum was largely unexpected; sterling experienced its greatest one-day loss since the introduction of free-floating exchange rates in the 1970s during the 24 hours following the vote, reflecting the adjustment of market expectations to the outcome of the referendum.

We find that regions that were *ex ante* more exposed to the threat of future trade barriers on professional services exports experienced a decline in online job postings after the referendum, relative to less exposed regions. This decline was economically meaningful: a one standard deviation increase in exposure to potential future barriers led to a 3.1% decrease in monthly job postings in our baseline specification, controlling for exposure to tariffs, the impact of the exchange rate depreciation, and immigration policy uncertainty. The effect was mainly felt by higher

[3]MFN terms are defined as follows: 'Under the WTO agreements, countries cannot normally discriminate between their trading partners. Grant someone a special favour (such as a lower customs duty rate for one of their products) and you have to do the same for all other WTO members.' More information can be found here: https://www.wto.org/english/thewto_e/whatis_e/tif_e/fact2_e.htm

skilled job adverts, with a particular impact on postings for executives, managers and professional occupations. When broken down by professional service sectors, the effect was concentrated primarily in financial services, information services, engineering and, to some extent, legal services. In contrast to services barriers, we find that the threat of future tariffs on goods does not seem to have impacted the posting of job adverts.

The financial services sector received a lot of attention during the Brexit negotiation period because of its importance to the UK economy. In 2018, financial services alone accounted for 6.9 percent of UK GDP and 11.7 percent of all UK services exports (ONS, 2018). The issue of 'passporting' rights, which allow financial businesses authorised in any EU Member State to operate freely across the EEA, was central to the Brexit discussions because the removal of these rights was perceived as an insurmountable barrier to many of the core business operations of financial firms based in the UK.[4] We therefore also separately explore the impact of this threat to the financial service sector on labour demand. We use the fact that the UK publishes data on the regional breakdown of exports of services, which is available for certain industry categories of which financial services is one. This allows us to exploit additional variation in the EU export intensity of the financial service sector across different UK regions. We find that local labour markets that were more reliant on EU export intensive financial services before the vote experienced a substantial decline in the posting of job adverts after the vote and this affected lower-skilled job adverts in addition to higher skilled ones.

After the Brexit referendum took place in June 2016, three and a half years passed before the UK officially left the European Union on the 31st January 2020. This period was characterised by a long and varying political process, including three general elections, two changes of Prime Minister, multiple failed votes in parliament and three extensions to the date when the UK would officially leave the EU. Businesses were left to interpret the political signals and adjust to regular changes in the implications of the proposed arrangements. In the second part of the paper, we therefore focus on the evolution of public perceptions and media

[4]For more information on passporting see https://www.bankofengland.co.uk/prudential-regulation/authorisations/passporting

coverage during this drawn out negotiation period.[5] We construct two measures of time-varying Brexit-induced trade-policy uncertainty. The first uses the intensity with which the UK's top ten newspapers were publishing articles mentioning both Brexit and uncertainty about future trading relationship with the EU. The second uses the Google search intensity in the UK for terms relating to both Brexit and trade policy.

Using these measures, as well as the Brexit Uncertainty Index (BUI) developed by Bloom et al. (2019), we show that the relative decline in online job adverts for regions more exposed to future professional services trade barriers was larger in months with heightened uncertainty over future trading policy arrangements with the EU. These results are not seen, however, when running a placebo test using the counterfactual of the postings that would have occurred had UK sectors followed the same trajectory as in the US, suggesting that the results do not just reflect sectoral trends.

We make a back-of-the-envelope calculation of aggregate effects, estimating that approximately one third of the total reduction in the posting of job adverts in the UK, relative to the pre-vote trend, can be explained by exposure to the threat of barriers on professional service exports. This is equivalent to around 1.5 million fewer job adverts being posted cumulatively between June 2016 and December 2019. Taken together, our results suggest that the threat of unravelling trade integration with a country's major trading partners can have substantial negative effects on labour markets through the reduced advertising of job openings. For the UK, it was the threat of barriers on trade in services, and particularly financial services, that mattered the most, rather than the threat of barriers on goods trade, despite the often greater focus on tariffs and goods supply chains during the discourse around the economic impacts of Brexit.

Our contribution to the literature consists of several elements. First, we contribute to the emerging literature looking at the immediate impacts of trade barriers and loss of market access. Amiti et al. (2019), Cavallo et al. (2019) and Fajgelbaum et al. (2019) study the consequences of the 2018 US trade war for prices, trade flows and welfare. In the context of Europe, Mayer et al. (2018) quantifies

[5] In Section 2.2 we provide a timeline of the key events during this period.

the trade-related welfare gains each country member has reaped from the European Union and the potential cost of reverting to a shallow regional agreement or WTO terms. We provide a new contribution by studying the impact of the potential loss of market access to the world's most integrated trading bloc on near real-time labour demand, proxied by online job adverts, and also by evaluating the relative labour market impact of barriers on both services and goods trade.

We also contribute to the literature on economic policy uncertainty. There is a substantial body of research showing that uncertainty affects investment, growth and employment. Rooted in the earlier work of Bernanke (1983) and Hassler (1996) who highlighted the importance of variations in uncertainty, Bloom (2009) develops a theoretical model whereby macro uncertainty shocks produce a rapid drop and rebound in aggregate output and employment as higher uncertainty causes firms to temporarily pause their investment and hiring.[6] We contribute both by focusing on trade-specific uncertainty and by investigating how the time-varying nature of trade policy uncertainty affects a real-time measure of labour demand.

We additionally contribute to the expanding literature looking specifically at trade policy uncertainty (e.g. Pierce and Schott (2016); Crowley et al. (2018a); and Handley and Limão (2017)). While these papers typically study the effect of trade policy uncertainty on trade, investment, and firm entry/exit, we consider how this feeds through to individuals through its impact on labour demand. Two recent papers that have studied the economic consequences of the uncertainty surrounding Brexit are Crowley et al. (2018b) and Graziano et al. (2018). The former paper looks at UK exporters' entry and exit from the EU market, using a difference-in-difference approach comparing firms that are differentially exposed to potentially high 'threat-point' tariffs before and after the referendum. The latter paper concentrates on uncertainty pre-referendum and its impact on the value of bilateral trade. The evidence to date has therefore tended to focus on tariffs and goods trade barriers. We provide new evidence on the impact of the services trade barriers posed by Brexit, in addition to goods trade barriers, and how both of these affected labour demand over the entire pre-referendum and negotiation period.

[6]For a review of the theoretical literature, see Dixit and Pindyck (1994). For a review of some of the empirical literature, see Baker et al. (2016).

Third, our paper contributes to a broader literature on the economic impacts of the Brexit referendum, providing new evidence on labour demand, another first-order consequence of the Brexit vote. Costa et al. (2019) show that that the sharp devaluation of the pound on the night of the Brexit referendum resulted in a negative effect on worker salaries and training post-referendum. We follow their method to control for the impact of exchange rate changes and show that the exchange rate depreciation also had an impact on labour demand that is distinct from that of trade policy uncertainty. Our findings also support the work of Bloom et al. (2019), who study the economic policy uncertainty caused by the Brexit vote and show that more productive, internationally exposed, firms have been more negatively impacted by the aggregate uncertainty caused by Brexit than less productive domestic firms; our results suggest that this effect could be operating through the threat of future barriers on services exports. Breinlich et al. (2019) show that the Brexit vote increased outward foreign direct investment from the UK into the other EU member states. Our results point to the possibility that this could be driven by firms trying to avoid potential future trade barriers.

Fourth, we add to a broad literature on understanding trade in services that stretches from examining the benefits of services liberalization (Arnold et al. (2011), Arnold et al. (2016), Barone and Cingano (2011), Beverelli et al. (2017), Breinlich et al. (2018)) to services offshoring (Liu and Trefler (2019), Crinò (2010), Eppinger (2019), Amiti and Wei (2005)) to the tradability of services (Jensen and Kletzer (2005), Gervais and Jensen (2019)). Regarding services trade and Brexit, Ebell (2016) and Mulabdic et al. (2017) highlight the particular importance of deep trade agreements for services trade and indicate that any deal outside the Single Market is likely to have a major impact on trade flows. We provide new evidence on how the threat of services trade barriers and unravelling deep trade agreements affects labour markets, demonstrating the important and often under-emphasised role that trade in services plays in advanced economies.

Finally, we add to a relatively new literature using real-time labour market data, such as online job adverts, to study labour markets (Hershbein and Kahn (2018); Deming and Noray (2018); Deming and Kahn (2017)), shedding light on the substantial decline in the posting of online job adverts in the UK after 2017.

This paper proceeds as follows: Section 2.2 provides background information on the referendum, Section 2.3 summarises the data sources used and construction of our exposure measures, Section 2.4 outlines the empirical strategy, Section 2.5 presents the results, Section 2.6 the robustness checks and Section 2.7 concludes.

2.2 Background on the Brexit Referendum

Table 2.1 provides a summary of the key events and dates relating to the referendum and the negotiation period.

2.2.1 The referendum

The UK electorate voted to leave the European Union on the 23rd June 2016 with a lead of 3.8%, against the positions held by the major political parties and to the surprise of many observers. The unexpected nature of this vote is evidenced by the betting markets that had placed the likelihood of a 'leave' outcome at around 30% for most of the preceding year, and the fact that the pound-dollar exchange rate fell by 8 percent in the 24 hours following the referendum (the sterling's biggest one-day loss since the introduction of free-floating exchange rates in the 1970s). Furthermore, the Ipsos Mori 'Issues Index' shows that less than 10% of those surveyed considered the EU, Europe, or Brexit as important issues in Britain during the ten years prior to the vote.[7] The unpredicted nature of the result is important for our analysis in that it permits us to compare the pre- and post-referendum periods in a difference-in-difference specification.

Understanding the result is complicated and multidimensional, but sovereignty was a key theme raised by the Vote Leave campaign as captured by the key slogan 'take back control'. The issue of sovereignty was of particular salience for issues of borders and immigration, UK trade policy, reallocating UK financial contributions to the EU to national concerns such as the NHS, and independence in making UK laws.[8] There is also evidence from Fetzer (2019) that the significant post-financial crisis austerity measures enacted by the UK

[7] https://www.ipsos.com/sites/default/files/ct/news/documents/2019-02/issues_index_january2019_v1_internal_use_only.pdf

[8] http://www.voteleavetakecontrol.org/why_vote_leave.html

Government led to increases in support for the UK Independence Party, which in turn strongly correlated with Leave support in the referendum.

2.2.2 Post-referendum uncertainty

The referendum was the start of an extended period of profound uncertainty about the future UK-EU relationship. The post-referendum timeline can be split into three key periods: the negotiation period following the vote but prior to leaving the EU; the transition period after leaving the EU; and the final deal. Our analysis covers the first period, starting after the referendum and ending on the 31st January 2019 with the UK leaving the EU. The primary objective during this period was to negotiate a legally binding 'withdrawal agreement' covering the terms of the transition period, and a non-legally binding 'political declaration' regarding aims for the final deal.

This period consisted of significant ups and downs including three general elections, two changes of Prime Minister, intense brinkmanship regarding leaving the EU without a deal, dramatically different competing visions for the future, and multiple failed votes in parliament. Table 2.1 presents the key events in the Brexit timeline. For each of the key policy areas affected by Brexit, it was unclear during the negotiation period which of many potential outcomes would be realised, and hence firms had to infer the probability of these outcomes from political signals. In a range of areas, to this day it remains unclear what regulations will be in place after the transition period. Consensus is now building that this has led to delayed economic responses to the referendum as firms and individuals waited for more clarity on specific issues and policies.

2.2.3 Possible future scenarios for trade policy

UK firms trading with the EU faced three main potential future trade arrangements upon leaving the EU, each with different regulatory barriers and tariff schedules: staying in the Single Market; leaving the Single Market and negotiating a comprehensive free trade deal; leaving with no deal.

The 'Single Market' outcome would ensure continued frictionless and tariff-free trade in goods and access to the European market for services. However, in

TABLE 2.1. BREXIT TIMELINE

Date	Event
23rd Jan 2013	David Cameron declares he is in favour of an EU referendum
14th Apr 2015	Launch of the Conservative Party Manifesto for the 2015 General Election committing to 'hold an in-out referendum on our membership of the EU before the end of 2017'
7th May 2015	Election of Cameron on Manifesto containing referendum promise
7th Sep 2015	European Union Referendum Act passed in parliament
20th Feb 2016	Date of referendum confirmed
23rd Jun 2016	EU Referendum
13th Jul 2016	Cameron steps down, Theresa May becomes Prime Minister
29th Mar 2017	Invocation of Article 50
8th June 2017	Snap General Election, Conservative party remains largest party but loses seats, relying on Democratic Unionist Party for majority in Parliament
20th Nov 2017	Michel Barnier states "The legal consequence of Brexit is that the UK providers lose their EU passport'
14th Nov 2018	The Withdrawal Agreement is agreed and published
25th Nov 2018	EU27 leaders endorse the withdrawal agreement and approve political declaration on future EU-UK relations
15th Jan 2019	First failed vote on withdrawal deal in UK Parliament
16th Jan 2019	UK Government wins vote of no confidence
12th Mar 2019	Second failed vote on withdrawal deal in UK Parliament
14th Mar 2019	Vote to request extension of Article 50 in UK Parliament (to 12th April if no deal agreed or 22nd May if deal agreed)
29th Mar 2019	Third failed vote on withdrawal deal in UK Parliament and originally planned leaving date
10th Apr 2019	The UK and EU27 agree to extend Article 50 until 31st October 2019
24th May 2019	May gives official notice of her resignation
24th Jun 2019	Boris Johnson elected Prime Minister by conservative party members
19th Oct 2019	New Brexit deal lost on amendment in the Commons, Prime Minister writes to European Council president to request another extension
12th Dec 2019	Johnson wins majority in the UK General Election
23th Jan 2020	European Union (Withdrawal Agreement) Act received Royal Assent
31st Jan 2020	UK leaves the European Union and enters transition period, due to run until end of 2020

Notes: This table shows the timeline of the events leading to the UK's exit from the European Union.
Sources: Commons Briefing papers CBP-7960, Nigel Walker, https://commonslibrary.parliament.uk/research-briefings/cbp-7960/.

this scenario the UK would not be able to freely negotiate its own trade agreements and would not have a say in future EU negotiations. Furthermore, the EU's position on the indivisibility of the four freedoms (of movement of goods, services, capital and people) meant that free movement of people would remain, and the UK would remain in the jurisdiction of the European Court of Justice. This option was ruled out at an early stage by Prime Minister Theresa May, with the EU then explicitly ruling out sector-specific arrangements that might have maintained existing benefits.[9] Single Market rules allow financial businesses authorised in any Member State to operate freely across the EEA, a system known as passporting. In November 2017, Michel Barnier, the EU's chief Brexit negotiator stated, 'The legal consequence of Brexit is that the UK financial service providers lose their EU passport''. This was the point at which it became relatively clear that the UK would not be able to maintain passporting rights.

The 'free trade deal' outcome would, in principle, allow the UK to maintain some preferential trade arrangements with the EU while retaining control over immigration, UK tariffs, and UK laws. Regarding goods, this meant no regulatory alignment but with the majority of goods being traded without being subject to tariffs. The challenges of the Irish border subsequently came to the fore, specifically the management of different regulatory regimes without having a physical (or 'hard') border and the potential to undermine key elements of the Good Friday Agreement that was designed to maintain peace in the region. The EU insisted on introducing a 'backstop' into the transition agreement, negotiated by Theresa May, to ensure that, in the absence of a mutually agreed solution to managing the border, at least Northern Ireland would remain in the customs union of the EU. However, this transition agreement failed three times to pass through parliament, leading Theresa May to give official notice of her resignation on the 24th May 2019.

For the financial services sector, the focus moved from passporting to 'equivalence'. Equivalence would allow UK and EU financial businesses to carry out specified activities across borders as long as the regulations that underpin these activities are deemed to achieve comparable outcomes. The process of agreeing and maintaining equivalence is neither straightforward nor as all-encompassing as

[9]See 8th January 2017 speech by Theresa May here: https://www.gov.uk/government/speeches/the-governments-negotiating-objectives-for-exiting-the-eu-pm-speech

passporting. For example, the approach can take time, depends on wider technical or political concerns, is limited to specified services, assumes continued close regulatory alignment between parties, and can be revoked by either side at any time. In light of this, there is evidence that by 2018 many financial institutions had already started to move their corporate headquarters, branches and staff away from the UK. New Financial (2019) reported in March 2019 that 275 financial institutions had moved or were moving at least part of their operations elsewhere in the EU.

The final possible outcome was a 'no deal' where, for services, the UK's exports would revert to being governed only by the terms set under the GATS, and for goods, the UK would export to the EU under the EU's WTO tariff schedule and pay MFN tariffs. Although widely considered as a very negative outcome for businesses, the option was first promoted by Theresa May when she argued that 'no deal is better than a bad deal' and later strengthened when Boris Johnson became Prime Minister with statements such as 'we are getting ready to come out on October the 31st. Come what may...Do or die. Come what may."[10] This rhetoric meant that, while it had previously been seen as very unlikely that the UK would end up without a deal, from mid 2019 it started to look like a serious possibility.

In general, the possibility of trade barriers for both goods and services became increasingly likely as the negotiation period progressed. For services trade, the decision not to remain in the Single Market in early 2017 was the first move towards a substantial change in regulation for UK services exporters. For services trade it was possible that even under a preferential trade agreement, important barriers of entry for UK firms in EU markets might remain, meaning that the outlook became bleak relatively early on in the process. For financial services particularly, the crunch point came at the end of 2017 when it became clear that passporting rights would not be maintained after leaving. For goods trade, on the other hand, increased tariffs were a possibility primarily under a no deal scenario, meaning that the risk became sizable later on. The average EU MFN tariff is also relatively low (with a trade weighted average of 3% in 2018), meaning that it was also less clear how much of a threat this posed to businesses in reality, with the exception of firms heavily reliant on global supply chains.[11]

[10]https://www.ft.com/content/366432ce-9741-11e9-8cfb-30c211dcd229
[11]https://www.wto.org/english/res_e/statis_e/daily_update_e/tariff_profiles/E28_E.pdf

TABLE 2.2. SUMMARY STATISTICS

Variables	Mean	Median	Min.	Max.	Std Dev.
Job postings:					
Monthly postings	2,442	668	1	206,743	10,225
Monthly postings - unskilled	775	265	1	48,718	2,449
Monthly postings - skilled	1,614	388	1	151,708	7,525
US monthly postings placebo	9,265	8,835	2,176	31,375	3,335
Exposure measures:					
Professional services exposure	0.127	0.0877	0.0323	0.548	0.0902
Financial services exposure	0.0335	0.021	0.00814	0.287	0.0346
Export tariff exposure	0.305	0.208	0.0204	4.2	0.412
Import tariff exposure	0.598	0.408	0.0423	5.52	0.732
Imported inputs tariff exposure	0.23	0.203	0.0671	0.775	0.11
Controls:					
Xrate impact via imports	0.0652	0.0652	0.0644	0.0665	0.000327
Xrate impact via exports	-0.0622	-0.0627	-0.0705	-0.0505	0.00347
EU national share	0.0465	0.0444	0.0182	0.144	0.0172
EU8 national worker share	0.0297	0.0274	0.0121	0.0624	0.0106
Uncertainty indices:					
Newspaper-based uncertainty index	4.33	4.12	2.16	7.15	1.45
Google-search-based uncertainty index	2.13	1.9	1.13	4.55	0.788
Brexit uncertainty index	44.3	39.2	34.8	58.3	7.79

Notes: This table displays summary statistics for all of the key variables used in the analysis. The dataset includes a total of 213 TTWAs, across 60 months (Jan 2015 - Dec 2019), totalling 12,780 observations. All uncertainty measures run from Sept 2016 - Dec 2019. The newspaper uncertainty index has been scaled by 100 so the mean of 4.33 represents 433 articles per month.

2.3 Data

2.3.1 Online job adverts

We use data collected by Burning Glass Technologies (BGT), a company that scrapes, parses and deduplicates online job postings on a daily basis and creates labour market analytics products using this data.[12] BGT identify that their postings are sourced from approximately 40,000 online job boards and company websites, resulting in a comprehensive dataset including nearly all job adverts posted online. For the UK there are around 60 million UK job adverts in their data over the

[12]Duplicates are recorded as a single posting in the first period in which the posting occurs.

period 2012-2019. In Appendix 2.A we provide a detailed discussion of BGT's data coverage relative to other sources of labour market data for the UK. We show that between 2012 and 2019 the number of job adverts included in BGT's data is approximately 86% of the total number of vacancies in the UK economy as reported by the ONS UK Vacancy Survey (UKVS), which conducts monthly surveys asking UK employers about their job vacancies.

The BGT data has the advantage of being very rich: BGT classify the job adverts along a range of dimensions, most importantly for us are the detailed classification by TTWA and Standard Occupational Classification (SOC) codes. However, online job adverts data inevitably only cover the subset of all vacancies advertised online and can suffer from issues of representation by industry, geography, occupation and time. Hershbein and Kahn (2018) provide a detailed analysis of the industry-occupation mix of vacancies in BGT relative to other detailed US data sources, such as JOLTS, and how this mix has changed over time. They find that the BGT postings are disproportionately concentrated in occupations and industries that typically require greater skill, but that the distributions are relatively stable across time and the aggregate and industry trends in the quantity of vacancies track official sources reasonably closely. Therefore, while online job adverts do not provide a complete picture of the entire labour market, they can provide a useful barometer of labour market demand, particularly when studying white collar occupations.

The postings cover 225 TTWAs across the UK, with 76% of total job postings being classified with a TTWA. The final dataset includes 213 TTWAs, due first to the exclusion of the 10 TTWAs in Northern Ireland from the BRES employment data, discussed below, and second to the exclusion of 5 small TTWAs which are not present in the BGT data. This leaves us with a final dataset with 213 TTWAs for 60 months, resulting in 12,780 observations. Table 2.2 displays summary statistics for monthly job postings and all other key variables used in the paper.

Over the period considered in our analysis, January 2015 to December 2019, a total of 31,208,288 postings are present in our dataset, translating into an average of 6,241,658 per year. Figure 2.2 displays the time series of UK monthly job postings over the period of our analysis. They follow an upwards trend before the referendum before flattening out and starting to decrease form mid-2018. We are

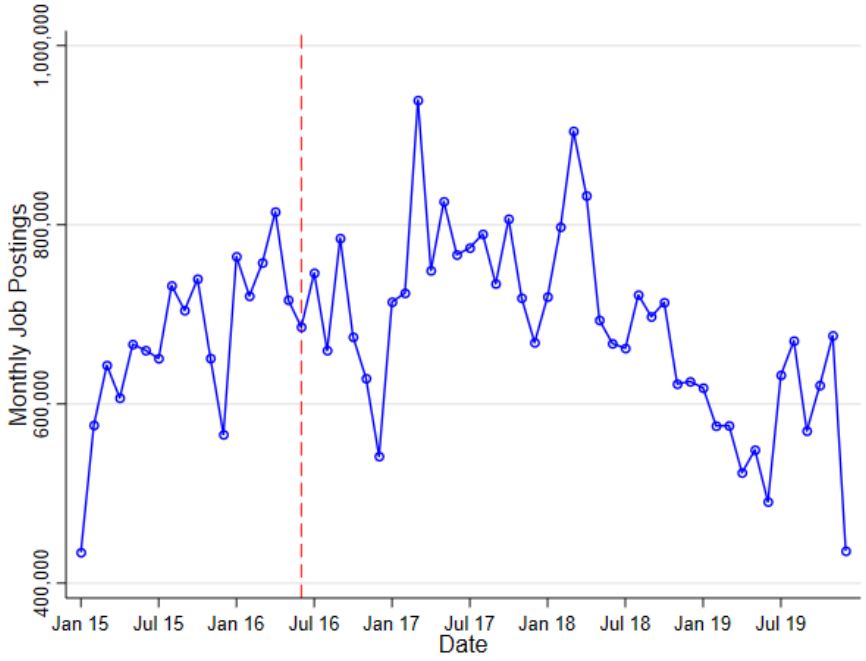


Figure 2.2: Monthly online job postings in the UK

Notes: Raw unsmoothed total monthly postings data. The vertical red dotted line identifies the month of the Brexit referendum.

interested in understanding to what extent concerns about a future relationship with the EU played into this decrease.

2.3.2 Local labour market exposure

Our analysis uses UK TTWAs as our statistical unit, which aim to reflect the geographic region where the population would generally commute to a larger town, city or conurbation for the purposes of employment.[13] The current criteria for defining TTWAs are that at least 75% of the area's resident workforce work in the area, at least 75% of the people who work in the area also live in the area, and the area must have an economically active population of at least 3,500. TTWAs range in population size from 6,800 to 8.4 million.

We use employment data from the UK Business Register and Employment Survey (BRES) for 2015 that contains a breakdown of employment by SIC4 industry

[13]TTWAs are defined by the ONS using census data for commuting between wards, based on the different locations of individuals' home and work addresses. See here: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/articles/traveltoworkareaanalysisingreatbritain/2016

within each TTWA in the UK.[14] BRES collects employment information from businesses across the whole of the Great Britain economy for each site that they operate, while the same information is collected separately for Northern Ireland. BRES surveys approximately 85,000 businesses, for 2015 the BRES data includes 28.5 million employees, 91% of the total UK labour force as estimated by the ONS.[15] As it is a business survey, the quality of the industry classifications is preferable to industry data from household surveys such as the Annual Population Survey, which we use for the immigration controls.

Our approach builds on the literature on local labour markets (e.g. Autor et al. (2014)). To calculate how exposed each local labour market is to potential future barriers to trade in professional services and goods, we first calculate a national sectoral exposure measure (defined in the following sections) before weighting this by the pre-sample sectoral employment composition for each TTWA using the BRES data discussed in the previous section.

$$\text{exposure measure}_r = \sum_{j \in r} \text{empl share}_{jr,2015} \times \text{exposure measure}_j \qquad (2.1)$$

where $\text{empl share}_{jr,2015}$ is employment in sector j and TTWA r as a fraction of total employment in TTWA r, and $\text{exposure measure}_j$ corresponds to the different trade barrier exposure measures presented in the following sections.

We ideally would like to use the employment composition before there was any possibility of Brexit. However, the sampling of BRES changed in 2015, substantially improving its coverage by including business units with a single Pay As You Earn code for which no Value Added Tax data are available. Prior to 2015, such units were excluded from the sampling frame and thus we choose to use 2015 data for the employment weights. We consider all employed individuals in a TTWA: an employee is defined as anyone aged 16 years or over that is paid directly from the payroll, in return for carrying out a full-time or part-time job or being on a training scheme. Employment includes employees plus the number of working owners who receive drawings or a share of the profits.

[14] Agricultural employment is excluded from the BRES, specifically SIC codes 0111-0150.

[15] https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/bulletins/uklabourmarket/2015-07-15

2.3.3 Professional services trade barrier exposure

Our aim is to quantify the threat of future trade barriers placed on UK exports of services if the UK were to leave the EU at the end of the transition period without a trade deal. If the UK and EU fail to reach an agreement, the UK will leave the EU's Single Market and will become a 'third country' to the EU for the purpose of services trade. Trade with the bloc will then follow the rules of the GATS of the WTO. In addition, UK businesses will face rules set by individual EU Member States.

In order to quantify the difference between the current barriers the UK faces on its services exports to the EU and the threat of what these barriers could revert to in a no-deal scenario, we make use of both the OECD STRI and the OECD Intra-EEA STRI.[16] The OECD STRI contains indices that measure MFN service trade restrictions for each country and sector, and does not take into account any specific concessions or preferential trade agreements. It was assembled by analysing laws and regulations in 34 OECD countries as well as Brazil, China, India, Indonesia, Russia, and South Africa. For each country and sector, five policy areas are considered: restrictions on foreign entry, restrictions on the movement of people, barriers to competition, regulatory transparency, and other discriminatory measures. The policy measures are grouped under the same five policy areas in all sectors, and are turned into an index using a scoring and weighing technique designed by the OECD. The indices take values from 0 to 1, with 1 indicating the highest Non-Tariff Measures (NTMs) (market completely closed to foreign services providers), and 0 meaning a fully liberalised sector.

Some examples of regulations included in the STRI for the Commercial Banking sector under the category 'restrictions on foreign entry' are: limiting foreign equity share in local banks, restricting cross border mergers, and requiring licenses. In the category of 'barriers to competition' some examples are: product level regulations, or having supervisory authorities that are not independent. For the Legal sector under the category 'restrictions on foreign entry' some examples are: whether commercial association is prohibited between locally and not locally

[16]OECD Services Trade Restrictiveness Index (STRI): https://qdd.oecd.org/subject.aspx?Subject=063bee63-475f-427c-8b50-c19bffa7392d. OECD Intra-EEA STRI: https://www.oecd-ilibrary.org/trade/intra-eea-stri-database_2aac6d21-en.

licensed lawyers, or whether acquisition and use of land and real estate by foreigners is restricted. For the category 'restrictions on the movement of people' some examples are: whether foreign professionals are required to take local exams, whether there are laws or regulations to establish a process for recognising qualifications gained abroad, or limitations on the duration of stay for intra-corporate transferees.

The intra-EEA STRI identifies and catalogues policy measures that restrict trade within the EEA for 25 OECD EU member countries. The information in the intra-EEA STRI is comparable with existing information in the STRI database but also covers EU law as well as national legislation. As a consequence, it follows that the resulting indices differ across EEA member countries, reflecting differences in national legislation. The OECD STRI could therefore be seen as a worst case scenario for the UK if it leaves the EU without a trade deal; trade restrictions will revert to MFN terms and barriers will be in line with those quantified in this index. The EEA-STRI, on the other hand, could be seen as the status quo of the current restrictions for the UK while it is in the EEA and covered by these EU laws.

We combine these two indices for the overlapping subset of 24 countries included in both and use them to evaluate the 'gap' between services trade restrictions placed on countries within the EEA and the MFN services trade restrictions they place on third countries.[17] We use the 2014 version of the STRI so that these restrictions were in place prior to discussion about Brexit. Of the 22 sectors, we exclude those relating to transport, logistics, construction and the arts (broadcasting, motion pictures and sound recording) such that we are left with an index that captures primarily professional service exports.[18] The sectors we include are then Accounting, Architecture, Commercial Banking, Computer services, Engineering, Insurance, Legal and Telecoms.

We use data from the Office for National Statistics (ONS) on UK exports of services by service type and country in 2015 to develop an EEA trade-weighted

[17]The 24 countries are: Austria, Belgium, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Netherlands, Norway, Poland, Portugal, Slovakia, Slovenia, Spain, Sweden

[18]We focus primarily on professional services due to both the feasibility of mapping between UK industries and STRI sectors, as well as the ease of interpreting the effect of trade restrictions on GATS Mode 1 services trade which comprises of services that are supplied across borders.

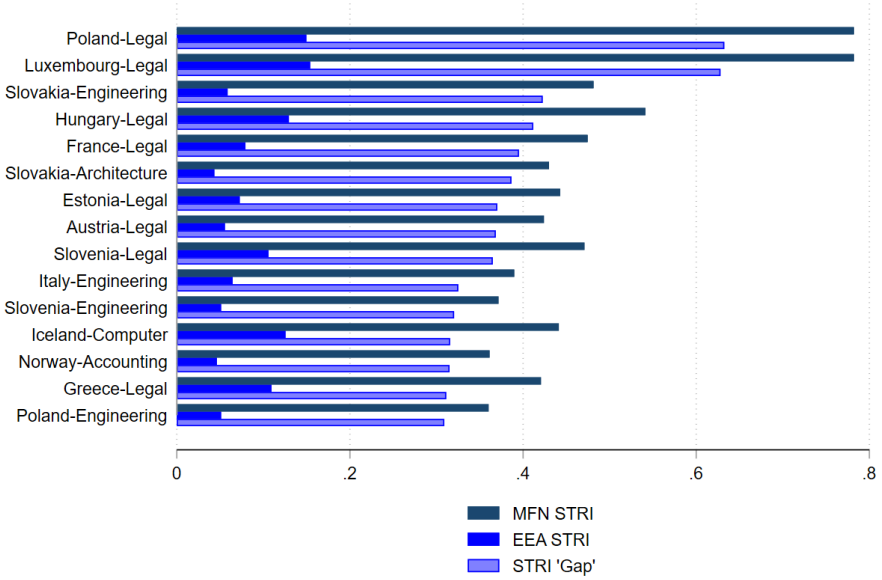


Figure 2.3: Country-sector pairs with highest STRI 'gap'

Notes: The figure presents the OECD's STRI ('MFN STRI') and Intra-EEA STRI ('EEA STRI') and the difference between them ('STRI gap') for the country-sector pairs with the highest gap. The STRI range from 0 to 1 with 1 representing the most restrictive. *Sources*: Services Trade Restrictiveness Index, OECD.

STRI 'gap' for each of the sectors we consider.[19] We map the ONS service types to the STRI sectors using the mapping described in Table 2.13 of Appendix 2.C.[20] These sectors can then be mapped to UK SIC 2007 codes, with this mapping also included in Table 2.13. We construct the sector level professional service barrier threat as follows:

$$\text{prof services exposure}_j = \frac{\text{Exports}_{j,2015}}{\text{L}_{j,2015}} \times \text{avg STRI gap}_{j,2014} \tag{2.2}$$

where the avg STRI $\text{gap}_{j,2014}$ is the trade weighted average difference between the 2014 MFN STRI and intra-EEA STRI for industry j across EEA countries, where the weights are EEA country shares in UK exports to the EEA in sector j. $\text{Exports}_{j,2015}$ are total exports to the EEA in sector j in 2015. Expressing exports in per worker

[19] We use the 2015 trade data because exports of services by detailed service type and country was not available prior to 2015

[20] We construct the trade-weighted sector level STRI gap using the UK-EEA export shares of only the subset of countries that are included in both of the STRIs and in the ONS export data, meaning this weighted-average is over 24 EEA countries.

terms follows the approach of Autor et al. (2013) and Blanchard et al. (2019). It allows us to capture tradability, i.e. the importance of exports for the given sector, as well as employment dependence on exports within the sector which maps directly to the decision of whether or not a firm decides to increase hiring and hence to post a job advert.

Figure 2.3 displays the top 15 country-sectors ranked by the difference between the MFN and intra-EEA STRIs. Legal services are the most commonly featured service type in this ranking, with the highest ranked differences being 'Legal services' in Poland, Luxembourg and Slovakia. For the trade-weighted STRI gap, the Legal sector also ranks highest, followed by Architecture, Accounting, Computer Services, Information Services, Engineering, Financial Services, Insurance then Telecoms with the lowest STRI gap.

We then construct measures of the professional services barrier exposure at the local labour market level as described in Section 2.3.2. Figure 2.4 presents a map of this measure by TTWA, with darker blue representing a greater professional services trade barrier exposure. There is an unsurprising concentration around the South East and London, combined with a number of very exposed areas in the North of England and Scotland. Figure 2.5 displays the 15 most exposed TTWAs: London ranks only 5th overall, with Edinburgh, Halifax, Trowbridge, Swindon and Skiption ranked as more exposued.

Financial services exposure

The UK is relatively unique in collecting detailed data on the regional breakdown of service exports. The ONS provides the value of service exports for aggregated service categories, of which Financial Services (henceforth FS) is one, originating from the 11 UK NUTS1 regions (for example North West, South West, Scotland), excluding Northern Ireland. This means that for FS specifically, we can can exploit additional regional variation in the export intensity of the FS sector which is not possible for the other professional services. Concentrating on one sector, however, comes at the cost of not being able to exploit variation across sectors of the STRI gap.

Within the group of professional services considered, FS is also of particular importance for the UK economy. FS alone accounted for 42% of the professional

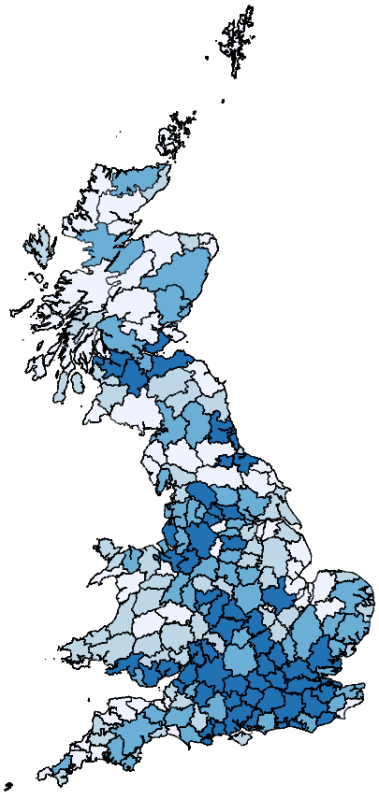


Figure 2.4: Exposure by Travel to Work Area: professional services trade barriers

Notes: This map displays the baseline employment-weighted professional services exposure for each TTWA. Darker colours represent more exposed areas.

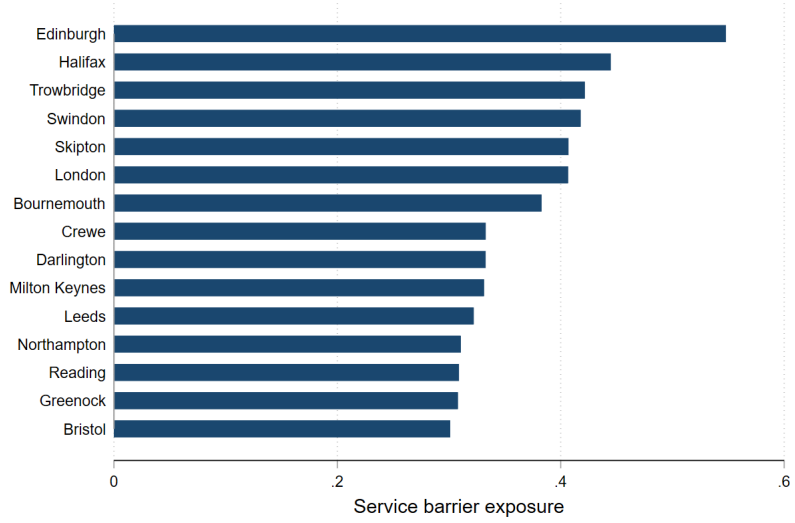


Figure 2.5: Most exposed Travel to Work Areas: professional services trade barriers

Notes: This chart presents the employment-weighted professional services exposure for the top 15 most exposed TTWAs. *Sources*: OECD, UN Comtrade.

services exports to the EEA in the ONS data used in 2015. The FS sector in the UK is also uniquely reliant on 'passporting' arrangements with the EU, which had a distinct and important role in the negotiations, as discussed above. The loss of passporting rights was considered one of the greatest potential consequences of Brexit as these rights were essential to many business functions of financial institutions based in the UK. In the second part of the analysis, we therefore focus solely on the FS sector and explore its impact separately. We construct a measure of exposure to future barriers on exports of FS as follows:

$$\text{Regional FS exposure}_{\text{NUTS1}} = \frac{\text{Regional FS Exports to EEA}_{\text{NUTS1,2015}}}{\text{L}_{\text{FS,NUTS1,2015}}} \times \text{avg STRI gap}_{\text{FS,2014}} \quad (2.3)$$

where Regional FS Exports to $\text{EEA}_{\text{NUTS1,2015}}$ are the FS exports of UK NUTS1 region $NUTS1$ to the EEA, trade-weighted avg STRI $\text{gap}_{FS,2014}$ is defined as above and $\text{L}_{FS,NUTS1,2015}$ is UK regional employment in FS. The FS exposure of TTWA r in NUTS1 region NUTS1 is then:

$$\text{FS exposure}_{r,NUTS1} = \text{FS empl share}_{r,2015} \times \text{Regional FS exposure}_{NUTS1} \quad (2.4)$$

where $\text{FS empl share}_{r,2015}$ is the Financial Service (FS) share of TTWA r employment in 2015. Figures 2.15 and 2.16 in Appendix 2.B display a map of this measure of financial service sector employment share by TTWA as well as the 15 most exposed TTWAs. The list is topped by London, Edinburgh, and Trowbridge. A number of the most exposed regions are home to banks or building societies, for example Skipton (4th most exposed) being home to Skipton Building Society, or Halifax (5th most exposed) being home to Halifax Building Society.

2.3.4 Tariff exposure

This section explains the construction of the sectoral export exposure to MFN tariffs. The tariffs used in the analysis are taken from World Integrated Trade Solution (WITS), and we select the applied MFN tariffs that the EU applies to imports coming from the rest of the world (excluding countries with which the EU has preferential trading arrangements). The data are aggregated at the 6-digit level of

the Harmonised System (HS6) and represent the simple average of tariffs across higher levels of disaggregation. We match these tariffs to UK exports to the EU-27 at the HS6 level and then, in order to calculate the average MFN tariff per sector, we match the combined dataset with 4-digit ISIC codes using crosswalks provided by the UN Statistics Division.[21]

From there we aggregate the tariffs to a per-worker sectoral measure as follows:

$$\text{tariff exposure}_{j,2014} = \frac{\text{Exports}_{j,2014}}{L_{j,2015}} \times \text{avg MFN tariff}_{j,2014} \tag{2.5}$$

where $\text{avg MFN tariff}_{j,2014}$ is the export-weighted average EU MFN ad valorem tariff across all HS6 products mapped to sector j, $\text{Exports}_{j,2014}$ are exports from the UK to the EU-27 of a particular HS6 product in 2014, and L_j is the total national employment in sector j in 2015. There are 419 sectors, of which we keep only the 147 manufacturing sectors i.e. those in SIC Division C. We use the 2014 tariffs for two reasons. First, we want to avoid the unlikely possibility that the EU might be strategically adjusting its MFN tariffs in anticipation of the possibility of Brexit. Second, as our analysis uses a trade-weighted tariff measure and we want to avoid the possibility of trade flows being affected by the referendum results.

We focus primarily on the effect of future MFN tariffs on UK exports, rather than UK import tariffs, for a few reasons. First, while the UK would not be able to control the tariffs placed upon its exports if it left the EU without a trade deal, it would be able to directly control its import tariffs. In addition, it was often suggested during the negotiation period that the UK would unilaterally place low, or zero, tariffs on imports if it were to leave without a deal.[22] We would therefore expect that the perceived risk of harm from future import tariffs would be substantially lower than the risk of harm from tariffs on UK exports, which were widely known to default to WTO terms if the UK left without a deal. This said, measures that take into account potential import tariffs are included in the robustness section.

[21] The relevant crosswalks can be fond here: unstats.un.org/unsd/classifications/econ/. We use CPC Ver 2.1 as an intermediate nomenclature between HS 2012 and ISIC Rev. 4. The employment data is provided at the UK SIC 2007 level (equivalent to NACE Rev. 2 up to the 4-digit level) but can be straightforwardly aggregated to the ISIC Rev. 4 level using a concordance from the same UN stats source.

[22] For example, see https://www.ft.com/content/d97854c2-2941-11e9-a5ab-ff8ef2b976c7

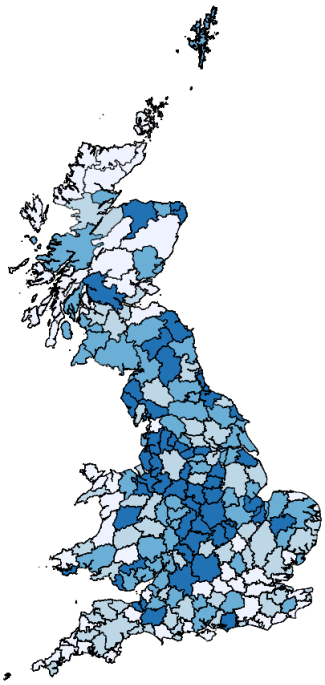


Figure 2.6: Exposure by Travel to Work Area: MFN tariffs

Notes: This map displays the baseline employment-weighted tariff exposure for each TTWA. Darker colours represent more exposed areas.

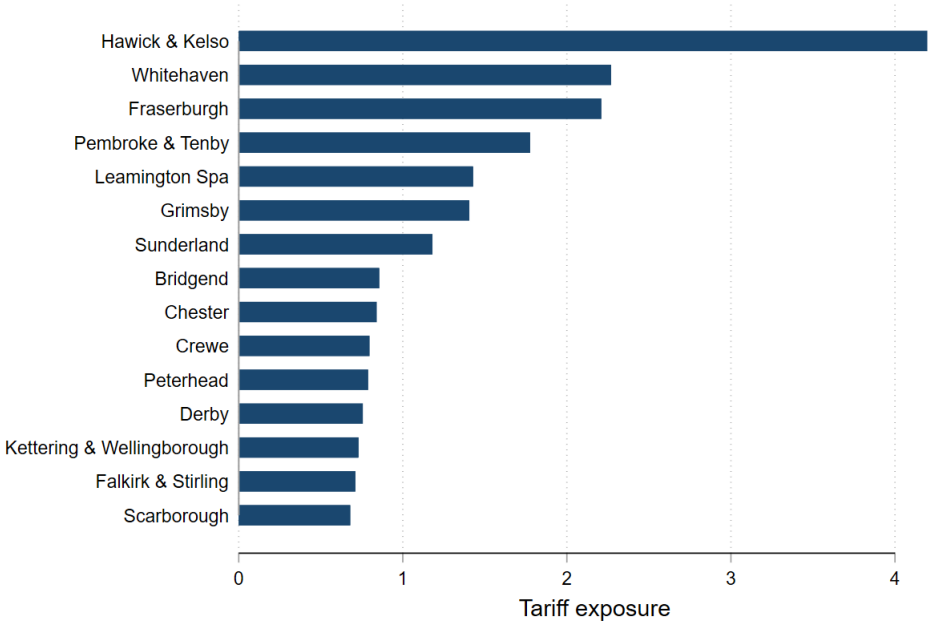


Figure 2.7: Most exposed Travel to Work Areas: MFN tariffs

Notes: This graph displays the MFN tariff threat for the 15 most exposed TTWAs. *Sources*: World Integrated Trade Solution (WITS), UN Comtrade.

TABLE 2.3. POSTING OCCUPATIONS AND SKILL LEVELS

SOC Code	Group title	Skill grouping
1	Managers, Directors and Senior Officials	3/4
2	Professional Occupations	4
3	Associate Professional and Technical Occupations	3
4	Administrative and Secretarial Occupations	2
5	Skilled Trades Occupations	3
6	Caring, Leisure and Other Service Occupations	2
7	Sales and Customer Service Occupations	2
8	Process, Plant and Machine Operatives	2
9	Elementary Occupations	1

Notes: Where a skill grouping of 4 is the highest skill level and 1 is the lowest as defined by the ONS. Our definition of high skill includes levels 3 and 4, and low skill includes 1 and 2. *Source*: ONS Standard Occupational Classification.

We then construct measures of the MFN tariff exposure at the local labour market level as described in Section 2.3.2. Figure 2.6 provides a map of the exposure measures and Figure 2.7 displays the top 15 TTWAs by exposure level. The most exposed TTWAs include Hawick & Kelso, Whitehaven, and Fraserburgh, all in the North of England or Scotland. The map shows a concentration in the Midlands with relatively little exposure in the South East and London, in contrast with the professional services exposure.

2.3.5 Classifying job adverts by occupation and skill

The job postings data provides information on the occupational classification of each of the postings at the 4 digit SOC level. Examples include 'Managers and proprietors in agriculture and horticulture' or 'Metal plate workers, and riveters'. These can be aggregated to nine 1-digit groups as presented in Table 2.3.

The ONS classifies the 2-digit sub-major groups of the SOC 2010 into four skill levels where 'skill level is defined with respect to the duration of training and/or work experience recognised in the field of employment concerned as being normally required in order to perform the activities related to a job in a competent and efficient manner'.[23]

We define the top two levels (3 and 4) as 'high skill' and the bottom two (1 and 2) as 'low skill'. Examples of high skill sub-major groups include 'Science,

[23]https://www.ons.gov.uk/methodology/classificationsandstandards/standardoccupationalclassificationsoc/soc2010/soc2010volume1structureanddescriptionsofunitgroups

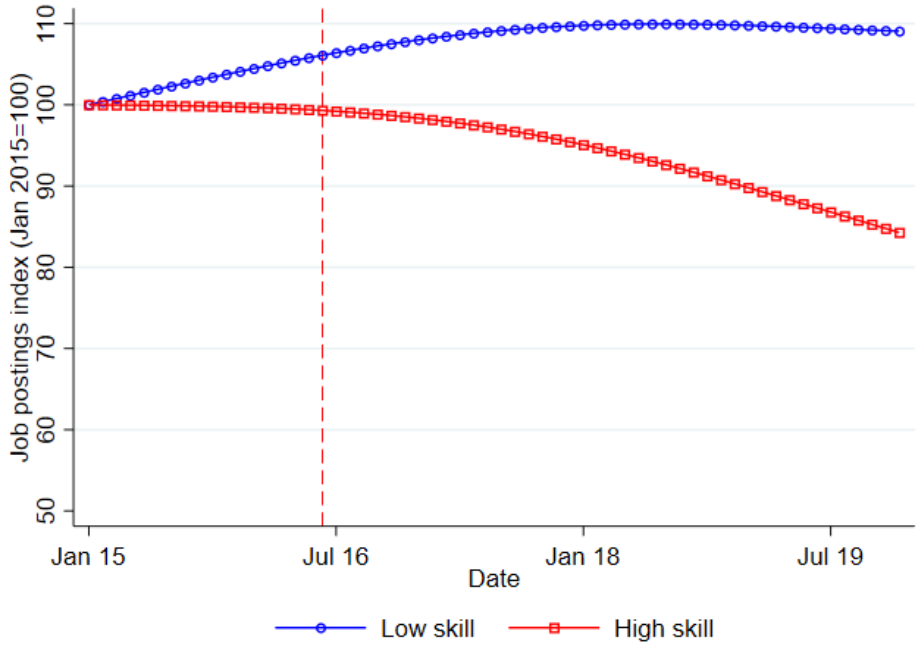


Figure 2.8: Skilled and unskilled job postings over time

Notes: Both trend lines are smoothed using the Hodrick-Prescott time-series filter, removing cyclical components, and normalised to 100 in January 2015 for comparison. The vertical red dotted line identifies the date of the Brexit referendum.

research, engineering and technology professionals' (level 4) and 'Business and public service associate professionals' (level 3). Examples of low skill sub-major groups include 'Administrative occupations' and 'Elementary trades, plant and storage related occupations'. Figure 2.8 displays the evolution of high skill and low skill job postings over time, we see a clear gap between the progression of the two types with high skill postings overall decreasing, and low skill postings slightly increasing over the period.

2.3.6 Measures of trade policy uncertainty

Our objective is to estimate the labour demand response to uncertainty surrounding the likelihood that UK firms will be subject to services barriers or MFN tariffs when exporting to the EU after Brexit. Although uncertainty is challenging to define and measure, we build on previous efforts in the literature and develop proxies which measure uncertainty using the time-varying prevalence of terms

mentioned in newspaper articles and google searches. We then compare the results using these indices to the measures of general Brexit-induced uncertainty in the BUI from the Decision Maker Panel (DMP) survey, constructed by Bloom et al. (2019), discussed more below.

Newspaper coverage

Our first approach uses UK newspaper coverage and builds closely upon the methodology in Baker et al. (2016)'s Economic Policy Uncertainty (EPU) Index and Ahir et al. (2018)'s World Trade Uncertainty (WTU) Index, but focuses specifically on the trade uncertainty caused by the Brexit referendum. We begin by combining the list of trade policy uncertainty related terms from the EPU and the WTU, which include words such as 'trade policy' or 'World Trade Organization'. We then remove all terms that would not be related to the trade uncertainty caused specifically by the Brexit referendum (for example, 'NAFTA' or 'Doha round'). This leaves us with a condensed list of 6 trade related terms: 'trade', 'tariffs', 'WTO', 'World Trade Organisation', 'trade policy', and 'trade agreement'. Given our additional focus on services trade, we also include two key services-restriction related words: 'passporting' and 'services agreement'. We then follow the WTU Index and search for articles that mention any of these terms with the words 'uncertain', 'uncertainty' or 'uncertainties'. While the EPU and WTU indices are interested in general trade policy uncertainty, in this paper we aim to isolate the trade policy uncertainty caused by the Brexit referendum. We therefore add an additional requirement for these terms to appear with the words 'Brexit', 'no deal', 'leave EU' or 'EU'.

Table 2.4 summarises these terms. We took a monthly count of any article including a term from Category A, a term from Category B, and a term from Category C. We search among the top 10 most popular UK newspapers by circulation: The Daily Mail, The Sun, The Mirror, The Express, The Times, The Telegraph, The Guardian, The Independent, The Daily Express and The Metro. Our data comes from Factiva, a news aggregator, and covers the period 2015-2019. This index is displayed in Figure 2.9 for the negotiation period, which is the focus of our analysis, normalised to 1 for September 2016 so that the three indices considered are

TABLE 2.4. UNCERTAINTY MEASURE INCLUDED TERMS

Category 1	Category 2	Category 3
brexit	uncertainty	trade
no deal	uncertain	tariffs
leave EU	uncertainties	passporting
EU		wto
		world trade organisation
		trade policy
		trade agreement
		services agreement

Notes: This table displays the terms used in the uncertainty measures. We counted any article including a term from Category 1, a term from Category 2 and a term from Category 3.

comparable for the period when the BUI is available, discussed below. Newspaper coverage peaked in November 2018, the month when the UK and EU finally agreed on the text of the draft withdrawal agreement and a summit was held where all EU27 nations endorsed the Brexit deal. Newspaper coverage remained high until March 2019, when the government put in a request to extend Article 50.

Google searches

Our second approach is to use an index of Google searches.[24] Google searches offer an alternative way to gauge the degree of public concern surrounding Brexit and future trade policy, through directly observing what people are searching for. Google Trends provides public data in terms of indices of the intensity of Google searches for specific terms conducted within a given region over time.[25] We construct an index of relative search intensity for the same terms as for the newspaper measure, but exclude the uncertainty-related terms. We assume that when individuals are uncertain about future trade arrangements they will not search for the word 'uncertain', whereas newspapers would report on uncertainty using these words.

[24] An alternative was to follow the approach in papers such as that by Graziano et al. (2018) and use betting markets to gauge uncertainty. However, betting markets tend to release contracts on narrowly defined questions over a limited period of time. Since we aim to measure uncertainty over the entire pre-and post-Brexit period, this type of measure was not feasible. Public polling was an additional option, but few polls asked the same question over time.

[25] https://trends.google.com/trends

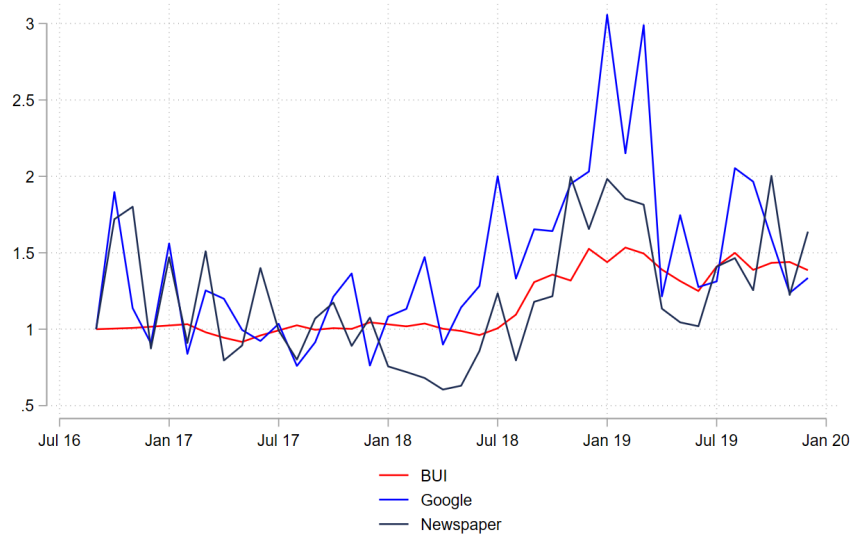


Figure 2.9: Brexit trade policy uncertainty measures

Notes: For illustrative purposes the measures are normalised to 1 on September 2016, non-normalised versions are used in regressions. The newspaper index is constructed using newspaper articles including key terms relating to Brexit, uncertainty and trade policy. The google index is constructed using google searches for key terms relating to Brexit and trade policy. The BUI originates from the UK DMP Survey.

Figure 2.9 displays the normalised google trends measure for the post-vote period alongside the two other measures. Similarly to the newspaper measure, this measure has a marked peak in late 2018 in a similar location to the peak of newspaper coverage, but the latter is much more muted. There is significant month-to-month variation across both measures.

BUI

In addition to the two uncertainty measures discussed above, we also compare our results using the BUI from the DMP survey.[26] The DMP is a panel survey of 8,000 firms, with around 3,000 responding in any given month. The BUI is defined as the share of firms which rate Brexit as one of the three highest drivers of uncertainty for their business.[27] Figure 2.9 compares the BUI with our two uncertainty measures from September 2016, when the BUI starts. All follow similar trends, although the BUI has a lower variance and fluctuates less than our measures.

[26]https://decisionmakerpanel.co.uk
[27]More information can be found in Bloom et al. (2019)

2.3.7 Control variables

The Brexit referendum also introduced other factors which may have affected the posting of online job adverts during this period. We focus on four key control variables capturing changes to the expected supply of EU or EU8 nationals working in the UK, and the sharp depreciation of the pound sterling after the referendum.[28] Below we discuss the construction of the immigration controls, and the exchange rate variables.

Accounting for immigration policy uncertainty

To measure the employment share of EU and EU8 nationals in a TTWA before the vote we use data from the Annual Population Survey (APS) in 2015. The APS is a continuous household survey covering the UK, with the aim of providing between-census estimates of key social and labour market variables at a local area level. The APS is not a stand-alone survey, but combines data from two waves of the main Labour Force Survey (LFS) with data collected on a local sample boost. The datasets comprise 12 months of survey data and are disseminated quarterly, with an achieved sample size of approximately 320,000 respondents. The APS is the most comprehensive source of data on employment by nationality of workers and is typically used for research on immigration in the UK.

The data provide a breakdown of the share of employment of EU and EU8 nationals in each UK NUTS1 region and SIC1 industry. We use data on the SIC1 employment composition of each TTWA in a given region to construct the employment share measures. A map of these measures is displayed in Figure 2.17 of Appendix 2.B.

Immigration was a central theme of the Brexit campaign and was one of the policy areas given priority during the negotiation period. The referendum result introduced substantial uncertainty surrounding freedom of movement of people between the UK and EU and the future ability of UK firms to employ EU nationals. We therefore introduce an additional control for a TTWA's share of employment of

[28]The EU8 countries are a group of eight of the 10 countries that joined the European Union during its 2004 enlargement. They are commonly grouped together separately from the other two states that joined in 2004, Cyprus and Malta, because of their relatively lower per capita income levels in comparison to the EU average. They are the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia and Slovenia.

EU nationals and EU8 nationals in the pre-referendum period, interacted with the post-referendum dummy. We hypothesise that firms relying on EU workers may take the referendum and risk of a no deal as a negative shock to their businesses and hence may reduce hiring until clarity of future arrangements is restored. The measure is defined as:

$$\text{EU national share}_{r,\text{NUTS1}} = \sum_{k \in r} \text{empl sh}_{kr,2015} \times \frac{\text{EU workers}_{k,\text{NUTS1}}}{\text{L}_{k,\text{NUTS1}}} \tag{2.6}$$

where r is the TTWA, NUTS1 is the NUTS1 region in which the TTWA is located, and k the SIC1 sector. Total workers, $\text{L}_{k,\text{NUTS1}}$, include any individual in employment, and EU workers include any EU national in UK employment. We construct an analogous measure for nationals of the EU8 member states, which include: Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia and Slovenia.

Accounting for the exchange rate depreciation

One of the most notable immediate impacts of the EU referendum was the large overnight depreciation of the pound with respect to the dollar and euro, the magnitude of which speaks to the unexpected nature of the referendum results. UK firms are likely to have been affected by this depreciation, both through increased cost of imported inputs, and through increased competitiveness of export products. The depreciation was also not equal with respect to different currencies, for example, the pound-dollar exchange rate fell by 8 percent overnight on June 23/24 while the pound-euro exchange rate fell by 6 percent. Since imports and exports differ in their source and destination countries, industries trading in different world markets faced a different effective sterling depreciation. The differential cost and revenue shocks from these country specific variations in the unexpected sterling depreciation therefore affected industries differently (Costa et al., 2019). If these sector-specific changes across time are correlated with the threat of services trade restrictions or MFN tariffs then we may be concerned that our key estimated impact is biased.

Following Costa et al. (2019), we include controls for the sector-specific (2-digit SIC) exposure to the exchange rate depreciation both in terms of exports and imported inputs. These controls are constructed as follows.[29] The intermediate import weighted exchange rate change (where the depreciation corresponds to a *positive* value i.e. more expensive inputs) is:

$$\hat{E}_o^M \equiv -\sum_i \sum_{s \neq uk} S_{sio} \hat{E}_s$$

where s indexes the source country for imported inputs, o the output industry, i the input industry, and M imports. $\hat{E}_s$ is the change in the exchange rate between the pound sterling and source country currencies, where a depreciation corresponds to $\hat{E}_s < 0$. S_{sio} is the share of intermediates of industry i from source country s in the factor expenditures of industry o.

The export weighted exchange rate change (where the depreciation corresponds to a *negative* value i.e. reduced prices of UK exports in the foreign currency) is:

$$\hat{E}_o^X \equiv \sum_{d \neq uk} S_{dxo} \hat{E}_d$$

where d indexes the destination country for exports, o the output industry, and x exports. $\hat{E}_d$ is the change in the exchange rate between the pound sterling and destination country currencies, where a depreciation corresponds to $\hat{E}_d < 0$. S_{dxo} is the share of destination d in exports of UK firms in output industry o.[30] We then construct measures of the exchange rate exposure at the local labour market level as described in Section 2.3.2. Figure 2.17 in Appendix 2.B provides maps of these measures by TTWA.

[29] We multiply both definitions by -1 to aid with interpretation in our particular context.

[30] This one day currency depreciation was not the only exchange rate fluctuation in the period studied. For example, the USD to GBP rate saw a maximum of 1.72 in July 2014 and a minimum of 1.23 in January 2017. However, given the lack of exogeneity of a concurrent exchange rate as a control, we are unable to control for other exchange rate fluctuations during the period, but control only for this one particularly sudden change.

2.4 Empirical strategy

2.4.1 Baseline specification

Our baseline specification estimates the impact of the post-referendum period on monthly online job postings in UK TTWAs as a function of a labour market's exposure to future potential trade barriers over the period of January 2015-December 2019. We estimate the following model:

$$\begin{aligned}\log(\text{postings}_{rt}) =& \beta_0 + \beta_1 \text{trade barrier threat}_r \times \text{post vote}_t + \text{X}_r \times \text{post vote}_t \\ &+ \gamma_t + \gamma_r + \epsilon_{rt} \end{aligned} \tag{2.7}$$

where postings_{rt} are the total number of online job adverts posted in month-year t and TTWA r, post vote_t is a dummy variable for the time period after the referendum, trade barrier threat_r is one of the measures of the exposure of TTWA r to future trade barriers between the UK and the EU, and X_r are the exchange rate and immigration controls.

We do not include the post-vote variable on its own as our specification includes month-year dummy variables. Equally we include TTWA fixed-effects to control for TTWA-specific, time-invariant factors. We are also interested in the impact of trade policy uncertainty for different skill and occupations groups and so we run the specification in (2.7) separately for each category of job postings, with these categories defined in the Data section.

2.4.2 Time-varying trade policy uncertainty

The outcome of the referendum introduced a large overnight changes in perceptions about future trade policy arrangements between the UK and EU. We proxy for this step change initially through our post vote_t dummy. However, the degree of uncertainty varied substantially in the period after the referendum as the negotiations continued with a significant lack of clarity regarding the different possible trade policy outcomes. We therefore also explore how the impact of exposure to professional service export barriers or future MFN tariffs varied during the post-vote period as a function of the proxies for Brexit-induced trade-policy uncertainty.

We consider the following specification:

$$\log(\text{postings}_{rt}) = \beta_0 + \beta_1 \text{trade barrier threat}_r \times \text{uncertainty}_t + \gamma_t + \gamma_r + \epsilon_{rt} \quad (2.8)$$

where uncertainty_t is one of the measures of uncertainty in month-year t. For this specification we focus on the post-referendum negotiation period, such that we evaluate differences in the impact of trade barriers within the period when the Brexit vote result had already been announced.

2.5 Results

2.5.1 Baseline post-vote results

We start with our baseline specification laid out in equation 2.7 before gradually adding controls. The results are presented in Table 2.5. Column (1) shows the coefficient on the interaction term between the post vote dummy variable and the TTWA trade weighted professional services exposure, and column (2) additionally includes the tariff exposure measure. Columns (3)-(5) subsequently add the controls described earlier in the paper. Across all columns, the professional services measure is negative and significant at the 1% level with the magnitude varying between -0.344 and -0.540. Taking column (4) as our preferred specification, a one standard deviation increase in exposure (0.09) leads to a 3.1% decrease in monthly postings. The tariff measure is close to zero and insignificant across all specifications.

Turning to controls, the export exchange rate measures are positive and significant across all specifications. This implies that TTWAs that were more exposed to a reduction in relative prices of UK exports abroad experienced a decrease in postings after the vote. This is surprising as we might have expected that the depreciation would have increased the competitiveness of exporting firms. We cannot ascertain what is driving this result but hypothesise that it could perhaps reflect the impact of the uncertainty caused by the exchange rate shock, or that the decline in relative prices for UK goods abroad was more than offset by more costly intermediate inputs. A one standard deviation increase in exchange rate export exposure (0.0034) is associated with an decrease in hiring of 3.1%. The import exchange rate measure is negative across all specifications. Its negative coefficient implies that

TABLE 2.5. BASELINE POST-VOTE RESULTS

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
post vote * prof service exposure	-0.538***	-0.540***	-0.351***	-0.346***	-0.344***
	(0.132)	(0.132)	(0.124)	(0.124)	(0.128)
post vote * tariff exposure		-0.008	0.001	0.003	0.001
		(0.033)	(0.035)	(0.035)	(0.035)
post vote * xrate impact via imports			-74.18	-81.09	-76.97
			(52.49)	(50.31)	(50.49)
post vote * xrate impact via exports			7.888*	9.084**	8.587**
			(4.194)	(4.065)	(4.277)
post vote * EU national share				0.639	
				(0.797)	
post vote * EU8 national share					0.556
					(1.341)
Observations	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.984	0.984	0.984
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES
Clustering	TTWA-YM	TTWA-YM	TTWA-YM	TTWA-YM	TTWA-YM

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable, with and without controls. All specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** p<0.01, ** p<0.05, * p<0.1.

exposure to relatively more expensive imported inputs due to the sterling depreciation reducing the posting of job adverts after the vote. However, the estimates just miss significance at the 10% level (p-value=0.11 for column (4)). Finally, the share of EU employees and EU8 employees are insignificant throughout.

2.5.2 Professional services impact over time

To evaluate how the professional services exposure affected postings throughout the negotiation period, we estimate a specification including interactions of the exposure measure with dummies for each quarter. Figure 2.10 plots the estimated coefficients. Quarters 1 & 2 from 2015 are excluded so all coefficients are relative to these quarters, and 95% confidence intervals are displayed vertically. We also control for the EU immigrant share and the post referendum exchange rate controls interacted with the post-vote dummy, month-year and TTWA fixed effects.

The pre-referendum coefficient estimates are close to zero and statistically insignificant, providing evidence that the professional services exposure was not affecting postings prior to the referendum. For several quarters after the

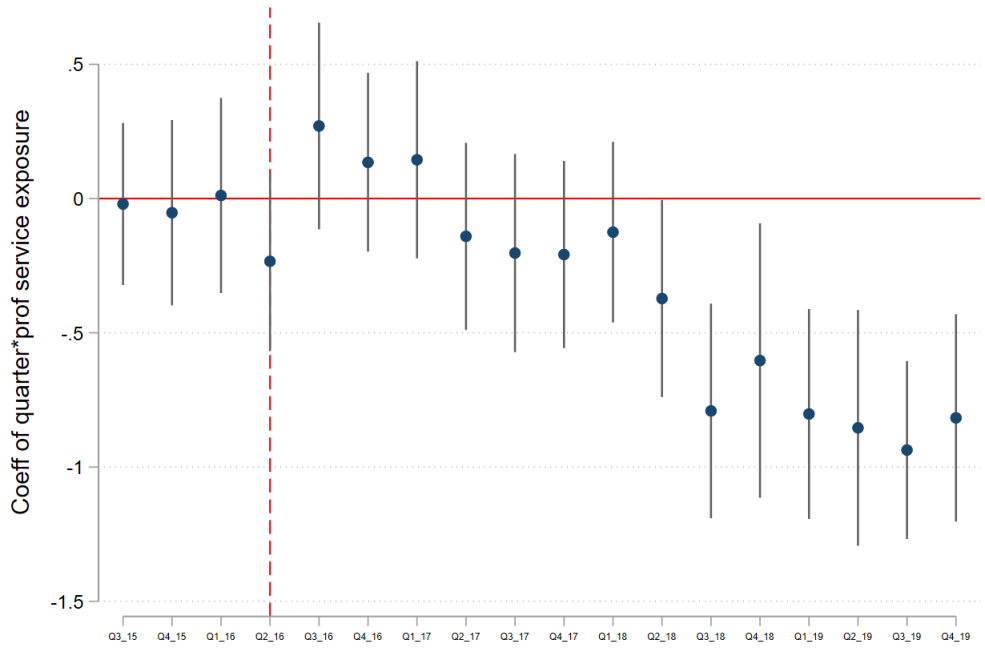


Figure 2.10: Impact of professional services exposure over time

Notes: This graph shows the coefficients from the regressions of the log of monthly job postings on the professional services barrier exposure measure at the TTWA level interacted with a dummy variable for each quarter from Q3 2015 to Q4 2019. Regressions also controlled for the EU immigrant share and the exchange rate measures interacted with the post vote dummy variables, TTWA fixed effects and month-year fixed effects. Coefficients are relative to the base period of Q1 2015 and Q2 2015. The dots represent the point estimates and the lines the 95% confidence intervals. The red line shows the quarter when the referendum occurred, Q2 2016.

referendum, the coefficients gradually become more and more negative, though remain insignificant. The first significant coefficient is in Q2 2018 followed by consistently negative and significant coefficients for the rest of the negotiation period. The delayed effect is consistent with the explanation that there was an initial period post-referendum where firms were trying to understand the future consequences of the vote. These results could imply that firms only responded by adjusting the posting of vacancies once it became clear that the likely outcome of negotiations was a significant separation from the Single Market.

2.5.3 Negotiation period uncertainty results

Our baseline results use a straightforward interaction of the professional services exposure with a post vote dummy. Although our specification controls for time-invariant TTWA-specific factors with the inclusion of TTWA fixed effects, and common time-varying factors are controlled for with year-month fixed effects, the approach might still be susceptible to the possibility that we are just capturing sector-specific variation over time.[31] In addition, 3.5 years passed between the referendum and the UK leaving the EU, a long time period in which a lot of political changes occurred and perceptions about the likely scenarios for Brexit will have changed substantially.

In order the address this concern, we interact our exposure measures with the three time varying uncertainty measures: the Google Brexit index, the Newspaper Brexit index, and the BUI. We focus on the period after the referendum starting in September 2016, the first date for which the BUI is available. The specifications now exclude the exchange rate controls as these pertain specifically to the pre-post referendum devaluation.

Table 2.6 presents the results for all three uncertainty measures. As before, the tariff exposure remains insignificant across 5 out of 6 specifications, confirming the results that businesses did not appear to be adjusting hiring as a response to a fear of no deal tariffs. For professional services exposure, the coefficients are negative and highly significant for both the Google and BUI measures, whereas the coefficients for the newspaper interaction are negative but weaker in significance. Taking the Google index as an example (column (1)), we see a coefficient of -0.23. For the mean exposure measure (0.127), an increase in uncertainty from the 25th to the 75th percentile (1.57 to 2.53) leads to a 2.8% decrease in monthly postings. The equivalent for the Newspaper measure is a decrease of 1.8%, and the BUI measure

[31] One such potential concern is the the EU service barriers or MFN tariffs tend to be high in declining industries because EU countries are trying to protect these industries from competition to slow down the process of job losses. If some UK regions in our analysis are dominated by such declining industries, we might mistakenly attribute their worsened job market performance to the Brexit shock. This concern is attenuated, however, by the fact that trade barriers are negotiated at the supranational level and the UK doesn't have direct control over the specific industries that get protected, although it may be able to achieve protection either through effective negotiation, or through shared interests with other EU states.

sees a decrease of 6.3%.[32]

We also now find negative and significant coefficients across all specifications for the EU and EU8 national shares, with the latter being larger in magnitude. Again taking column (1), for the mean value of the EU national shares (0.047) an increase from 25th to 75th percentile of the Google index decreases monthly postings by 8.5%. The equivalent for the EU8 share is a reduction of 8.3%. This result is consistent with the hypothesis that regions with a greater employment share of EU (or EU8) nationals before the vote were negatively affected by the threat of changes to immigration policy and so reduced the posting of vacancies in periods when there was heightened uncertainty relating to Brexit.

Table 2.15 in Appendix 2.C also presents additional specifications, which provide supporting evidence for the specific importance of monthly variation in these uncertainty measures above and beyond sectoral trends. All specifications make use of the full sample period and instead include a post vote * uncertainty measure * professional services exposure interaction term.[33] The results for the Google and BUI indices remain robust and the newspaper uncertainty index now generates highly significant negative results. We then iteratively control for either the post vote or quarter interaction terms, demonstrating that the results for Google and BUI remain robust to these specifications.

2.5.4 US placebo test

To provide further evidence that our results are not capturing global sectoral trends affecting the professional services sectors at the same time as our uncertainty measures peaked, we run a placebo regression using US postings mapped to UK TTWAs as the dependent variable. The aim of this exercise is to evaluate how job adverts would have responded, had the UK followed similar sectoral trends in job adverts to those that occurred in the US.

[32]The larger effect for BUI may be linked to the fact that it covers a broader definition of uncertainty than the trade specific measures.

[33]For the BUI, given it starts in September 2016, we extrapolate backwards using the September value for July and August 2016.

TABLE 2.6. UNCERTAINTY MEASURE RESULTS

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)	(6)
	Google		Newspaper		BUI	
measure*prof service exposure	-0.231***	-0.268***	-0.066*	-0.083*	-0.034***	-0.039***
	(0.058)	(0.060)	(0.039)	(0.042)	(0.007)	(0.007)
measure *tariff exposure	-0.007	-0.001	0.008	0.011**	-0.000	0.000
	(0.011)	(0.011)	(0.005)	(0.005)	(0.002)	(0.001)
measure*EU national share	-1.907***		-0.826***		-0.264***	
	(0.349)		(0.239)		(0.042)	
measure*EU8 national share		-2.900***		-1.249***		-0.413***
		(0.524)		(0.398)		(0.062)
Observations	8,520	8,520	8,520	8,520	8,520	8,520
Adjusted R-squared	0.987	0.987	0.987	0.987	0.987	0.987
TTWA FE	YES	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with varying uncertainty measures. The post vote period from September 2016 to December 2019 is considered. All specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** p<0.01, ** p<0.05, * p<0.1.

To construct a placebo at the UK TTWA level, we take the sector level time varying postings data from the US and weight it by the 2015 BRES pre-vote employment shares as follows:

$$\text{US postings}_{rt} = \sum_{j \in r} \text{empl share}_{jr,2015} \times \text{US postings}_{jt}$$

We first map the 6-digit North American Industry Classification System (NAICS) code, the US sector classification, to the 4 digit SIC codes, the UK equivalent.[34] Although there is a chance that US postings were also affected by the Brexit vote, we assume that this impact in the US would be far more muted than in the UK. If the professional services sectors that are threatened by high export barriers in the EU were globally on a downward trend relative to other sectors, and if our results were being driven by these trends occurring at similar times to when there was greater uncertainty, then we would expect to find similar results to our baseline in this US placebo regression.

Table 2.7 repeats the analysis in Table 2.6 with uncertainty interactions, now replacing UK postings by our US postings placebo. We find that none of our

[34]We use a crosswalk from the US Census Bureau (https://www.census.gov/eos/www/naics/concordances/concordances.html)

TABLE 2.7. US PLACEBO TEST

Dep variable: log US postings	(1)	(2)	(3)	(4)	(5)	(6)
	Google		Newspaper		BUI	
measure*service barrier exposure	0.002	0.002	-0.001	-0.001	-0.000	-0.000
	(0.017)	(0.016)	(0.007)	(0.007)	(0.002)	(0.002)
measure *tariff exposure	-0.001	-0.001	-0.000*	-0.000*	-0.000	-0.000
	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
measure*EU national share	-0.037		0.002		-0.011*	
	(0.052)		(0.033)		(0.007)	
measure*EU8 national share		-0.058		-0.006		-0.019*
		(0.078)		(0.048)		(0.011)
Observations	8,520	8,520	8,520	8,520	8,520	8,520
Adjusted R-squared	0.994	0.994	0.994	0.994	0.994	0.994
TTWA FE	YES	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA based on US sectoral trends on the average trade barrier exposure measures at the TTWA level interacted with varying uncertainty measures. The period September 2016 to December 2019 is considered. All specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

variables are statistically significant, even at the 10% level, hence supporting the claim that the changes we see in the UK are specifically caused by the threat of trade barriers and were not patterns also experienced in the US at these times.

2.5.5 Impact on different skill groups

Table 2.8 displays our baseline regressions run separately for high skilled and low skilled job adverts, the classification of which is described in the Data section. Panel (a) shows the effect for high skilled job adverts, while panel (b) shows the effect for low skilled job adverts. Relative to the baseline coefficients, we find that the high skilled professional service exposure coefficients are all strongly statistically significant and larger in magnitude, whereas the low skilled coefficients are mostly not significant. Taking column (4) high skill results, we see that a one standard deviation increase in exposure leads to a 3.9% decrease in postings (compared to 3.1% for aggregate postings). This is intuitive, given professional services typically involve higher skilled jobs, particularly those that are involved in exporting.

We also find a negative and weakly significant coefficient for the import exchange rate measure for high skilled postings, where a one standard deviation

increase leads to a 2.9% decrease in high skill postings. This is consistent with the explanation that firms that were more exposed to increases in the cost of their imported inputs reduced hiring for high skilled positions. The same effect is found for increasingly competitive exports, where a one standard deviation change in exchange rate leads to a 3.1% increase. Again, this is less straightforward to interpret but may be due to a step increase in uncertainty. EU shares are insignificant across both postings measures.

2.5.6 Breakdown by subsector

Our measure of professional services exposure comprises of variation across 9 subcategories. In order to better understand what is driving the change in job postings, Table 2.9 shows the results with the professional services exposure broken down into its constituent parts.

The overall professional services results are primarily driven by 3 out of the 9 subcategories: financial services, information services, and engineering services. Given the importance of financial services in these regressions and in the UK economy more broadly, as well as its prevalence in the discussions surrounding a future Brexit deal, in Section 2.5.8 we focus solely on this sector to better understand its specific impact on hiring.

2.5.7 Aggregate effects

In order to provide a rough idea of what these results for professional services may have implied for the full sample, we carry out a back-of-the-envelope counterfactual exercise on our baseline results. For each TTWA, we consider what the predicted values for monthly postings after the vote would have been if they had the exposure score of the 10th percentile TTWA in terms of professional services exposure. This exercise relies on two simplifying assumptions. First, the total effects on postings equal the sum of the direct effects on postings at each TTWA. Second, TTWAs below the 10th percentile in terms of professional services expo-

TABLE 2.8. IMPACT BY SKILL GROUP

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
	Panel (a) High skilled				
post vote * prof service exposure	-0.637***	-0.635***	-0.433***	-0.428***	-0.429***
	(0.132)	(0.132)	(0.137)	(0.138)	(0.141)
post vote * tariff exposure		0.009	0.018	0.020	0.018
		(0.032)	(0.034)	(0.034)	(0.034)
post vote * xrate impact via imports			-81.38*	-88.19*	-82.67*
			(47.58)	(46.30)	(46.23)
post vote * xrate impact via exports			7.753*	8.936**	8.077*
			(4.108)	(3.968)	(4.230)
post vote * EU national share				0.632	
				(0.811)	
post vote * EU8 national share					0.258
					(1.379)
Observations	12,773	12,773	12,773	12,773	12,773
Adjusted R-squared	0.982	0.982	0.982	0.982	0.982
	Panel (b) Low skilled				
post vote * prof service exposure	-0.280**	-0.283**	-0.180	-0.173	-0.164
	(0.127)	(0.127)	(0.119)	(0.117)	(0.121)
post vote * tariff exposure		-0.017	-0.008	-0.004	-0.007
		(0.037)	(0.039)	(0.038)	(0.039)
post vote * xrate impact via imports			-32.19	-41.68	-37.99
			(53.76)	(51.41)	(51.55)
post vote * xrate impact via exports			7.478	9.123*	8.933*
			(4.746)	(4.651)	(4.823)
post vote * EU national share				0.875	
				(0.764)	
post vote * EU8 national share					1.152
					(1.321)
Observations	12,766	12,766	12,766	12,766	12,766
Adjusted R-squared	0.976	0.976	0.976	0.976	0.976
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings by skill level in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable, with and without controls. All specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 2.9. PROFESSIONAL SERVICE RESULTS BREAKDOWN BY SECTOR

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
post vote * finance barrier exposure	**-0.499*** **	**-0.492*** **	**-0.461*** **	**-0.454*** **	**-0.471*** **
	(0.142)	(0.142)	(0.149)	(0.152)	(0.153)
post vote * engineering barrier exposure	**-6.284*** **	**-6.208*** **	**-5.731** **	**-5.501** **	**-6.068** **
	(2.238)	(2.256)	(2.490)	(2.636)	(2.722)
post vote * info services barrier exposure	**-3.815** **	**-3.858** **	**-3.740** **	**-3.804** **	**-3.696** **
	(1.778)	(1.772)	(1.785)	(1.775)	(1.759)
post vote * legal barrier exposure	-3.544	-3.630	**-4.154***	**-4.157***	**-4.221***
	(2.231)	(2.231)	(2.248)	(2.242)	(2.308)
post vote * insurance barrier exposure	1.136*	1.138*	1.040	1.032	1.046
	(0.589)	(0.588)	(0.651)	(0.655)	(0.651)
post vote * telecoms barrier exposure	-0.468	-0.545	-0.589	-0.544	-0.668
	(0.942)	(0.931)	(0.863)	(0.865)	(0.868)
post vote * computer barrier exposure	-0.800	-0.788	-0.065	-0.098	-0.038
	(0.783)	(0.780)	(0.890)	(0.894)	(0.884)
post vote * accounting barrier exposure	-1.904	-1.965	-1.774	-1.825	-1.740
	(2.794)	(2.811)	(2.970)	(2.969)	(2.963)
post vote * architecture barrier exposure	-5.654	-9.794	1.665	2.906	-1.725
	(57.43)	(58.78)	(56.38)	(56.63)	(57.87)
Observations	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.984	0.984	0.984
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable, with and without controls. All specifications include controls for the sterling depreciation's impact on imports and exports, as well as the EU national share, and TTWA and month-year fixed effects. Standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

sure did not change their posting of job ads after the vote in response to the threat of future service export barriers.

We define the counterfactual post-vote monthly postings of TTWA r had they had an exposure score in the τth percentile as:

$$\text{log postings}^{\tau}_{rt} = \text{E}\Big[\text{log postings}_{rt} | \text{exposure}_r = \text{exposure}^{\tau}_r\Big] \tag{2.9}$$

$$= \widehat{\text{log postings}}_{rt} + \hat{\beta}\left[\text{exposure}^{\tau}_r - \text{exposure}_r\right] \tag{2.10}$$

where $\hat{\beta}$ and $\widehat{postings}_{rt}$ are the fitted values from our baseline regression column (4). Then the total monthly loss of postings due to the threat of trade barriers is:

$$\text{Total monthly loss} = \sum_{\text{r: exposure}_r > \text{exposure}^{\tau}_r} \Big[\exp(\text{log postings}^{\tau}_{rt}) - \exp(\widehat{\text{log postings}}_{rt})\Big] \tag{2.11}$$

Taking τ as the 10th percentile, this exercise implies that after the vote the mean monthly decline in the number of job postings for a TTWA relative to this counterfactual was 164. Across all TTWAs each month there were 34,980 fewer job adverts posted than this counterfactual benchmark. The mean monthly number of job adverts posted prior to the vote across all TTWAs was 535,397, so this full sample effect implies a 6.5% loss in monthly job adverts on average after the vote for the full sample. Summing up over the whole post-vote period, this implies a cumulative loss of approximately 1.5 million job postings.

If we extrapolate the growth rate in total annual job postings in all TTWAs between 2012-2015 to the post-vote period, comparing this to the observed number of job adverts posted implies that there were 4.1 million fewer job adverts cumulatively in the post-vote period than had the UK continued to follow its pre-vote trajectory. The ratio of the total counterfactual lost postings to the actual reduction relative to the pre-vote trend is 0.36, implying that the threat of professional services trade barriers can explain around a third of the relative decline in job postings.

2.5.8 Financial services results

We now rerun the baseline regressions in Table 2.5 replacing professional services exposure by regional financial services exposure. As before, we find negative and

TABLE 2.10. FINANCIAL SERVICES POST-VOTE BASELINE

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
post vote * financial services exposure	-1.201***	-1.201***	-0.817***	-0.868***	-0.826***
	(0.296)	(0.295)	(0.279)	(0.254)	(0.269)
post vote * tariff exposure		-0.003	0.004	0.007	0.004
		(0.033)	(0.035)	(0.035)	(0.036)
post vote * xrate impact via imports			-98.19*	-105.3**	-100.6**
			(50.48)	(48.40)	(49.15)
post vote * xrate impact via exports			8.203*	9.857**	9.344**
			(4.238)	(4.070)	(4.260)
post vote * EU national share				0.914	
				(0.804)	
post vote * EU8 national share					0.940
					(1.310)
Observations	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.984	0.984	0.984
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable, with and without controls. All specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** p<0.01, ** p<0.05, * p<0.1.

significant results across all specifications with coefficients varying from -1.201 to -0.826. Taking column (4) again, we see that a one standard deviation increase in financial services exposure (0.0345) leads to a 3% decrease in monthly postings. We again see significant coefficients for the exchange rate impacts, where a one standard deviation increase in import exposure leads to a decrease in postings of 3.4% and the same measure for export exposure leads to an decrease of 3.4%.

Figure 2.11 displays the coefficients from the regression of log-postings on the financial services exposure measure interacted with quarter dummies. The results are very similar to the professional services case shown in Figure 2.10 with the impact being delayed until Q2 2018, and no effect showing prior to the referendum. Table 2.11 presents the financial services results broken down by skill level. As for professional services, we find that the results are stronger for high skill postings with a one standard deviation increase in exposure leading to a 3.6% decrease in postings. The exchange rate impacts are significant as for professional services with small changes in magnitudes.

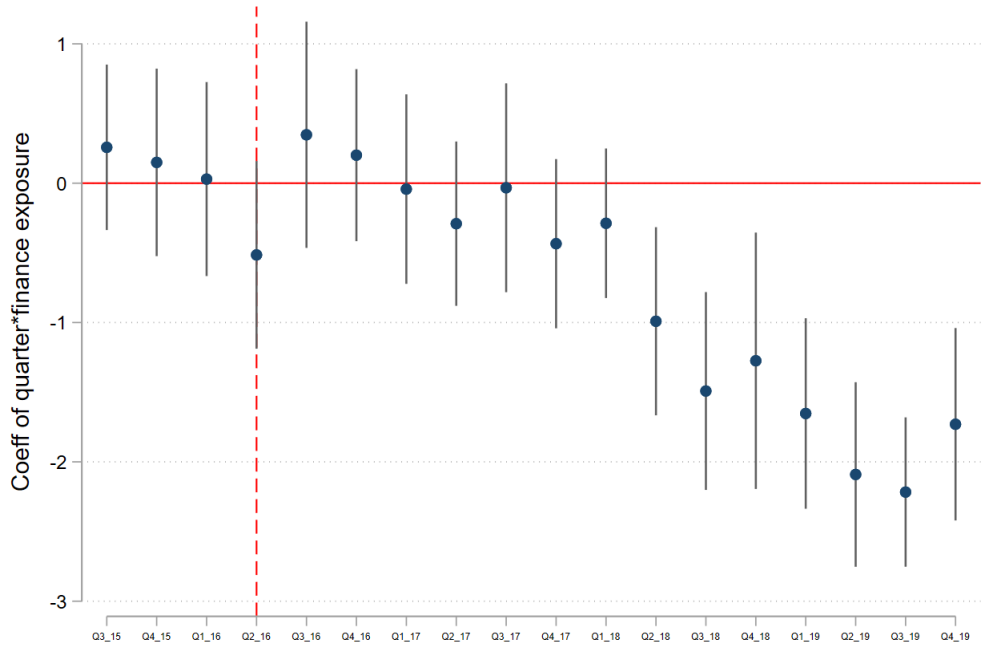


Figure 2.11: Impact of financial services exposure over time

Notes: This graph shows the coefficients from the regressions of the log of monthly job postings on the financial services exposure measure interacted with a dummy variable for each quarter from Q3 2015 to Q4 2019. Coefficients are relative to the base period of Q1 2015 and Q2 2015. The regressions also controlled for the EU immigrant share and the exchange rate measures interacted with the post vote dummy variables, TTWA fixed effects and month-year fixed effects, and standard errors are two-way clustered at the TTWA and month-year level. The dots represent the point estimates and the lines the 95% confidence intervals. The red line shows the quarter when the referendum occurred, Q2 2016.

TABLE 2.11. FINANCIAL SERVICES IMPACT BY SKILL GROUP

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
	Panel (a) High skill				
post vote * FS exposure	-1.414***	-1.414***	-0.989***	-1.043***	-0.996***
	(0.322)	(0.323)	(0.317)	(0.286)	(0.308)
post vote * tariff exposure		0.015	0.021	0.025	0.021
		(0.032)	(0.034)	(0.035)	(0.035)
post vote * xrate impact via imports			-111.6**	-119.0***	-113.5**
			(44.88)	(43.29)	(43.78)
post vote * xrate impact via exports			8.154*	9.900**	9.046**
			(4.159)	(3.979)	(4.213)
post vote * EU national share				0.964	
				(0.822)	
post vote * EU8 national share					0.734
					(1.346)
Observations	12,773	12,773	12,773	12,773	12,773
Adjusted R-squared	0.982	0.982	0.982	0.982	0.982
	Panel (b) Low skill				
post vote * FS exposure	-0.662**	-0.662**	-0.455*	-0.512**	-0.468*
	(0.275)	(0.274)	(0.262)	(0.237)	(0.249)
post vote * tariff exposure		-0.014	-0.0063	-0.002	-0.006
		(0.037)	(0.039)	(0.038)	(0.039)
post vote * xrate impact via imports			-43.27	-51.25	-46.79
			(51.63)	(49.68)	(50.31)
post vote * xrate impact via exports			7.613	9.489**	9.251*
			(4.750)	(4.618)	(4.775)
post vote * EU national share				1.032	
				(0.775)	
post vote * EU8 national share					1.345
					(1.300)
Observations	12,766	12,766	12,766	12,766	12,766
Adjusted R-squared	0.976	0.976	0.976	0.976	0.976
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings by skill level in each TTWA on the average trade barrier exposure measures interacted with the post vote dummy variable, with and without controls. All specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** p<0.01, ** p<0.05, * p<0.1.

To understand which specific occupation in this less intuitive lower skilled group were affected, we further explore the impact on job adverts by UK SOC group. Table 2.12 displays the same regressions run separately for job adverts for each of the nine high level occupation categories, the classification of which is described in Table 2.3 of the Data section. The coefficient on the FS exposure measure is strongly significant and negative for 5 out of the 9 occupations: *Managers, Directors and Senior Officials*; *Professional Occupations*; *Associate Professional and Technical Occupations*; *Skilled Trades Occupations*; and *Sales and Customer Service Occupations*. The coefficients range from -1.105 to -0.630 for *Managers, Directors and Senior Officials* and *Sales and Customer Service Occupations* respectively. A one standard deviation increase in financial services exposure then decreases postings for *Managers, Directors and Senior Officials* by 3.8%, the next highest coefficient is *Professional Occupations* for which the associated decrease is 3.6%.[35]

The *Skilled Trades Occupations* group includes a range of occupations, such as electricians, IT engineers, construction and building trades and even food preparation and hospitality trades. In general, these results suggest that local labour markets that were more exposed to barriers on FS exports reduced the posting of not only high-skilled jobs in the FS industry, but across a range of occupations, suggesting that the impact of FS spilled over to other types of jobs.

2.6 Robustness checks

2.6.1 Excluding London

A primary concern for these results could be that this impact was primarily driven by London, widely considered as the UK's professional services hub. First, Figures 2.4 and 2.5 show that the professional services exposure is actually not as concentrated in London as is often presumed. Additionally, Table 2.17 in Appendix 2.C repeats the analysis in Table 2.5 but excludes the TTWA of London. The resulting coefficients are extremely close to the baseline case with the same levels of significance, allowing us to conclude that London is not driving the results. FS exposure

[35]In Table 2.16 of Appendix 2.C, we also provide this analysis for the professional services exposure measure, finding that only the first three occupation groups were affected for professional services.

TABLE 2.12. FINANCIAL SERVICES IMPACT BY OCCUPATION

Dep var: log SOC postings	1. Managers, Directors and Senior Officials	2. Professional Occupations	3. Associate Professional and Technical Occupations
post vote * FS exposure	-1.105***	-1.030***	-0.878***
	(0.327)	(0.305)	(0.326)
post vote * tariff exposure	-0.018	0.059	0.006
	(0.031)	(0.047)	(0.038)
post vote * xrate impact via imports	-90.90*	-180.1***	-105.4**
	(49.14)	(47.15)	(52.70)
post vote * xrate impact via exports	1.741	13.64***	11.22**
	(4.354)	(4.158)	(5.558)
post vote * EU national share	1.133	0.559	1.309
	(0.880)	(0.833)	(0.957)
	5. Skilled Trades Occupations	7. Sales and Customer Service Occupations	4. Administrative and Secretarial Occupations
post vote * FS exposure	-1.004***	-0.630**	-0.365
	(0.313)	(0.291)	(0.292)
post vote * tariff exposure	0.009	-0.034	0.009
	(0.023)	(0.024)	(0.034)
post vote * xrate impact via imports	47.36	-43.75	56.40
	(61.94)	(52.67)	(56.14)
post vote * xrate impact via exports	2.340	2.284	9.635*
	(4.652)	(4.399)	(4.875)
post vote * EU national share	2.526**	1.176	1.309
	(0.969)	(0.918)	(0.951)
	6. Caring, Leisure and Other Service Occupations	8. Process, Plant and Machine Operatives	9. Elementary Occupations
post vote * FS exposure	-0.056	-0.612*	-0.474
	(0.277)	(0.327)	(0.318)
post vote * tariff exposure	0.035	-0.036	0.035
	(0.034)	(0.036)	(0.058)
post vote * xrate impact via imports	-68.27	51.62	23.61
	(51.68)	(61.31)	(59.99)
post vote * xrate impact via exports	18.98***	4.639	5.929
	(4.274)	(6.313)	(5.315)
post vote * EU national share	0.332	2.395**	1.051
	(0.840)	(0.987)	(0.945)
Observations	12,780	12,780	12,780
TTWA FE	YES	YES	YES
Month-Year FE	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings by occupation in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable. All specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** p<0.01, ** p<0.05, * p<0.1.

is also spread quite widely across the country, as is shown in Figure 2.15 of Appendix 2.B. Excluding London also does not have much impact on the significance or magnitude of the baseline results for FS exposure.

2.6.2 Including shares

A further concern is that our results are driven mainly by the shares in our exposure measure, rather than variation in the STRI or tariff component across sectors (the 'shift'). We therefore control for these shares for both the professional services measure and the tariff measure. The shares are constructed by a combination of the TTWA-sector employment weightings, and the sectoral exports to employment ratio, i.e excluding the trade weighted STRI or tariff component of the exposure measures.

Table 2.18 in Appendix 2.C repeats the baseline analysis in Table 2.5 but includes these shares. The professional services measure remains significant across all specifications, while the associated share is positive and significant at the 10% level in 4 out of the 5 specifications. This suggests that while professional services sectors that exported a lot to the EU were increasing their job postings, those which were more exposed to higher potential trade barriers experienced a relative decline in postings. To interpret the magnitude of the professional services coefficient in this specification, we fix the professional services share at its mean (0.82) and consider a one standard deviation increase in the sectoral STRI exposure (0.058) which leads to a 16% relative decline in monthly postings. Both the tariff exposure, and the related share, remain insignificant across specifications, although the exposure is now consistently negative across specifications.

2.6.3 Alternative tariff exposure measures

There are a range of ways in which tariff exposure has been measured in the literature. We additionally consider a number of other formulations of the exposure measure to check robustness. Table 2.19 in Appendix 2.C presents the results for four different exposure measures. Columns (1) and (2) use the 'output weighted' exposure. Instead of weighting by sectoral exports per worker, this mea-

sure weights by the fraction of EU exports in total sectoral output.[36] The exposure used in Columns (3) and (4) removes this weighting altogether, leaving a simple export weighted average tariff normalised by neither total employment nor total output. Columns (5) and (6) first take $ln(1 + \tau)$ before applying weightings but otherwise use the same construction as the baseline measure. Finally columns (7) and (8) adapt the baseline by replacing 4-digit national exports by 2-digit regional exports in the weighting.

The results suggest that these different formulations of the exposure measure do not significantly change the conclusions from the baseline and provide confidence in our result that firms did not appear to be adjusting hiring significantly in response to the threat of potential future MFN tariffs.

2.6.4 Pre-vote placebo

In order to further address the concern that our baseline pre-post results are driven by time varying unobservables that are not absorbed by month-year fixed effects, we conduct a further placebo test. Here we limit the data to only 2015 and split the sample into a pre- and post-period with six months in each. In other words, we introduce a break in mid-2015 instead of using the actual timing of the referendum, which took place in June 2016. We then interact a post-June 2015 dummy with the professional services exposure. Table 2.23 in Appendix 2.C presents the results of this alternative placebo test using pre-referendum postings. The resulting coefficients are not significant across all specifications, suggesting that this impact was not observed prior to the referendum, but driven primarily by the result of the vote.

2.6.5 Goods non-tariff barriers

Although tariffs are the most conspicuous trade barrier for goods, it could be argued that after decades of global negotiations they are reduced to a level which is no longer material for firms in most countries. Perhaps more concerning are the non-tariff barriers (NTBs) to trade in goods, which may remain even if the UK were to sign a free trade agreement with the EU. In order to investigate this possibility

[36] Due to data limitations, this weighting is at the 2-digit ISIC level and is sourced from the UK's Office of National Statistics input-ouput tables for 2014

we create a new measure of regional exposure to NTBs on exports of goods to the EU.

The measure is constructed is an identical way to the baseline tariff exposure measure but with tariffs replaced by a product-level NTB exposure. We use data from the World Integrated Trade Solution which provides a list of MFN non-tariff barriers at the 8-digit HS level. We first count how many barriers each HS8 product would be exposed to under MFN terms and then take a simple average at the HS6 level to match with our Comtrade export data. This measure is then weighted by each product's share of exports to the EU before replacing the avg MFN tariff$_{j,2014}$ in equation 2.5.

Table 2.20 presents the results, NTB exposure is not statistically significant in any specification, generally mirroring the results found for the tariff exposure that the threat of barriers on exports of goods does not seem to have had an impact on the posting of job adverts at the local labour market level.

2.6.6 Specific tariffs and quotas

Alongside ad valorem tariffs, calculated as a percentage of the good's value, there are also relevant non-ad valorem tariffs. Two key cases are specific tariffs, computed on the physical quantity of the good being imported, and tariff-rate quotas, made up of a low or zero tariff rate on an initial increment of imports (the within-quota quantity) and a very high tariff rate on imports entering above that initial amount. Following Crowley et al. (2018a) we use data from WTO Tariff Analysis Online to calculate similar tariff threat measures for the EU MFN specific tariffs and quotas that would be applied under a no deal, specifically creating trade weighted averages of dummies indicating whether a particular product has a specific tariff, a quota, or either of the two.[37] We then include three weighted average measures for each of these three trade barriers as a robustness check on our results.

Table 2.21 presents the baseline results with these measures included. Relative to our baseline result (column (1)) which only considers ad valorem tariffs, the professional services and ad valorem tariff results change very little with the former remaining strongly negative and significant, and the latter insignificant. The

[37]WTO Tariff Analysis Online: https://www.wto.org/english/tratop_e/tariffs_e/tariff_data_e.htm

specific tariff and quota measures are not significant in any of the specifications confirming the previous results that there is evidence for the negative impact of the threat of future barriers on goods trade.

2.6.7 UK import tariffs

Although most of the discussion around tariffs centered around the potential impact of EU tariffs on UK exports, there was also some uncertainty concerning possible UK tariffs on imports from the EU under a no deal scenario. We therefore further evaluate whether the threat of these tariffs had any impact on the posting of job adverts. We consider two potential channels.

Exposure to reduced competition through import tariffs

One channel through which import tariffs could affect UK businesses is by increasing the price of imports that compete with UK products, hence rendering UK firms more competitive in the domestic market. We can create a similar exposure measure as for exports but with weightings based on UK imports rather than exports. Although it was unclear during the negotiation period what the UK tariff schedule would look like under a no deal scenario, we use the EU MFN tariffs as a plausible default option for the analysis. This is reinforced by the fact that the MFN principal ensures that the UK could not unilaterally lower its tariffs with respect to the EU without doing the same thing for imports from third countries, unless part of a comprehensive free trade agreement (not present under 'no deal' by definition). The potential sectoral future tariff protection exposure is:

$$\text{imp protection}_{j,2014} = \frac{\text{Imports}_{j,2014}}{L_{j,2015}} \times \text{avg MFN tariff}_{j,2014}^{(Imp)} \tag{2.12}$$

where $\text{avg MFN tariff}_{j,2014}^{(Imp)}$ is the import-weighted average EU MFN ad valorem tariff across all products mapped to sector j, $L_j, 2015$ is the national employment in sector j (4-digit ISIC), and $\text{Imports}_{j,2014}$ is UK imports from the EU in 2014. The TTWA level exposure is then:

$$\text{imp protection}_r = \sum_{j \in r} \text{empl share}_{jr,2015} \times \text{imp protection}_{j,2014} \tag{2.13}$$

where empl share$_{jr,2015}$ is industry j's share of TTWA r employment.

Exposure to increased cost of imported inputs

Alongside potential protection of UK industries, tariffs may have the additional negative impact of increasing the cost of inputs. If industries typically import inputs which are either not produced by UK firms, or are produced at a higher price, then the imposition of tariffs on these products would increase costs and potentially reduce production. We calculate this exposure by taking the import protection measure from above and, using UK input-output tables, weighting it by the share this 'input' industry makes up in all of the 'output' industry's imported inputs.[38] Specifically, the measure is calculated as follows:

$$\text{intinputs threat}_{k,2014} = \frac{1}{L_k} \sum_j S_{j,k} \sum_{p \in j} \text{Imports}_p \times \text{MFN tariff}_{p,2014} \tag{2.14}$$

$$\text{intinputs threat}_r = \sum_{k \in r} \text{empl share}_{rk,2015} \times \text{intinputs threat}_{k,2014} \tag{2.15}$$

where k is the output sector, l is the input sector (both at the 2-digit SIC level), and $S_{l,k}$ is the imported inputs from l as a share of total imported inputs by k.

The results for both import tariff measures are presented in Table 2.22. As was the case for the export exposure measure, we find no effect of potential import tariffs on job adverts throughout the negotiation period. Given that we would expect stronger effects for the export exposure measure, as the default no deal tariff schedule was more clearly defined, it is perhaps not all that surprising that we find no effect for these other measures.

2.7 Conclusion

This paper uses data on the near universe of UK online job postings from January 2015 to December 2019 to analyse how the threat of future trade barriers on UK exports of both services and goods to the EU affected labour demand throughout the

[38] UK input-output tables: www.ons.gov.uk/economy/nationalaccounts/supplyandusetables/datasets/ukinputoutputanalyticaltablesdetailed

Brexit negotiation period. We exploit *ex ante* variation in EEA country-sector MFN service trade restrictions relative to restrictions applicable to non-EEA countries to develop measures of local labour market exposure to the threat of no-deal scenario export barriers for professional services exports. We construct an analogous measure for goods exports based on EU MFN tariffs.

We find that the threat of trade barriers on professional services exports to the EU had a strong negative effect on the posting of online job adverts after the Brexit referendum. A one standard deviation increase in local labour market exposure to future professional service export barriers decreased monthly postings by 3.1% on average in the post-referendum period. A back-of-the-envelope counterfactual calculation suggests that the aggregate impact of this exposure across all TTWAs implies a cumulative 1.5 million postings lost between July 2016 and December 2019, equivalent to around 1.2 million per year (for reference, the size of the UK labour force is approximately 31 million). We provide evidence that this negative impact kicked in around Q2 2018 and remained negative for the rest of 2018 and 2019. It is possible that this delayed effect is due to firms taking time to reassess and understand the implications of trade barriers for their operations, or that the likelihood of the UK leaving the Single Market and facing stringent restrictions on service exports looked low straight after the vote but subsequently increased substantially after the first year of negotiations.

In contrast, the threat of future MFN tariffs on goods exports to the EU does not seem to have affected online job adverts during the negotiation period. The same also holds for import tariffs and NTBs on goods exports. We provide a few ideas as to why this may have been the case. While MFN tariffs in some cases can be large, at the HS6 level 63% of product lines would have tariffs below 5%. These tariffs would cut away at profitability but may not be game-changing in a way that some restrictions on services exports, such as revoking passporting rights, or requiring commercial presence to export engineering or legal services to certain EU countries that were previously major trade partners, could be. One possibility is that these tariffs were simply not a large enough threat to have caused firms to stop hiring prior to being levied. Another explanation could be that there was substantial heterogeneity in how different manufacturing industries were affected. When

broken down by subsectors, we discovered that some sectors, such as auto manufacturing, did decrease hiring, but others, such as chemicals, increased hiring. Perhaps only industries hit by tariffs multiple times through global value chains faced substantial enough repercussions to pause hiring. Finally, it could just be that manufacturing represents a relatively small share of UK employment (9.6% in 2015) and so impacts on the manufacturing sector may not be large enough to show up when studying local labour market outcomes.[39]

Given the Brexit negotiation period lasted 3.5 years, we also explore the role of trade policy uncertainty in affecting job postings during this period. We develop two time-varying measures of Brexit-related trade policy uncertainty during this period based on newspaper coverage and Google searches. Additionally, we make use of the Brexit Uncertainty Index from Bloom et al. (2019). Using these indices we show that uncertainty was a key driver of the reduction in postings. We interact our trade barrier exposure measures with these uncertainty measures during the negotiation period and find that the impact of the threat of barriers on exports of professional services was greater in months with heightened uncertainty. Taking the Google index, for the mean exposure measure, an increase in uncertainty during the negotiation period from the 25th to the 75th percentile lead to a 2.8% decrease in postings. Using US job postings data to construct a counterfactual for each UK TTWA, we show that such a pattern would not have occurred in the UK, had the UK followed US sectoral trends in online job adverts.

We find that the threat of professional services export barriers only affected the posting of higher skilled job adverts. When broken down into sectoral components, it was the threat of barriers on financial services, engineering, and information service exports that had the biggest impact. We explore the impact of barriers to financial services exports specifically, finding strong negative impacts of exposure to EU-export intensive financial service exports on the posting of job adverts, that affected not only higher-skilled but also lower-skilled job adverts. We show that these results all hold when additionally controlling for other key channels through which the vote could have affected labour demand, including the

[39] UK total employment data by sector can be found here:
https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/datasets/employmentbyindustryemp13/current

sharp exchange rate depreciation in the 24 hours following the referendum and immigration policy uncertainty.

We conclude that the threat of future trade barriers on the export of professional services caused by the Brexit referendum had a far greater effect on labour demand than the threat of tariffs on goods exports, despite the greater focus of the Brexit discourse on manufacturing and tariffs. This impact on job adverts was substantial and indicates that a retreat from trade integration, and particularly deep integration, can have important consequences for labour markets.

Acknowledgements We thank Costa et al. (2019) for sharing their constructed exchange rate variables, and Crowley et al. (2018b) for sharing their NTB variables. We are grateful to Bledi Taska and Layla O'Kane at BGT for their help and support facilitating access to the job postings data. We thank Abi Adams, Nick Bloom, Steve Bond, Swati Dhingra, Kevin O'Rourke, Ferdinand Rauch, Tom Sampson, Daniel Trefler, Gabriel Ulyssea, Chris Woodruff. David Weinstein, and seminar participants at Oxford, UCL, LSE, FREIT and Villars for insightful comments and feedback.

Appendices of Chapter 2

2.A Comparison of BGT data with other sources

Online job adverts are one indicator of labour demand. The number of job adverts posted online encompasses only a subset of labour demand because a) not all vacancies will be advertised publicly and b) not all publicly advertised vacancies will be advertised online. We would hence expect the BGT job advert data to differ from the total number of actual vacancies in sectors, regions and job types where employers are less likely to publicly advertise vacancies or to advertise them online. In the UK, it is not a legal requirement to publicly advertise a job vacancy. However, there is an obligation for employers not to discriminate against employees or potential employees and an employer could face legal action if it is believed that a job has not been fairly advertised. Consequently, it is very common for firms to have company policies that require all open positions to be publicly advertised.

We would therefore generally expect that BGT would cover a high proportion of all vacancies, particularly for large firms. However, if one job advert is posted online for a number of openings at once, for example in the case of graduate schemes, then we would expect that BGT would underestimate the true number of vacancies.

In this section, we first evaluate the relationship between BGT online job adverts and another comprehensive measure of labour demand: the estimated number of vacancies from the UK Vacancy Survey.[40]

1. UK Vacancy Survey (UKVS)

The UKVS is a statutory, monthly survey of businesses conducted by the ONS.[41] The survey asks a single question: how many job vacancies did a business have in total (on a specified date) for which they were actively seeking recruits from outside their organisation? Results from the survey cover all sectors of the UK economy and all industries, with the exception of employment agencies and agriculture, forestry and fishing.

[40] The UK Vacancy Survey is limited to national quarterly figures, is based on a significant amount of estimation, and doesn't provide a breakdown by occupations, hence the reason we do not use this data source for the main analysis of the paper.

[41] A summary can be found here: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/methodologies/vacancysurveyqmi

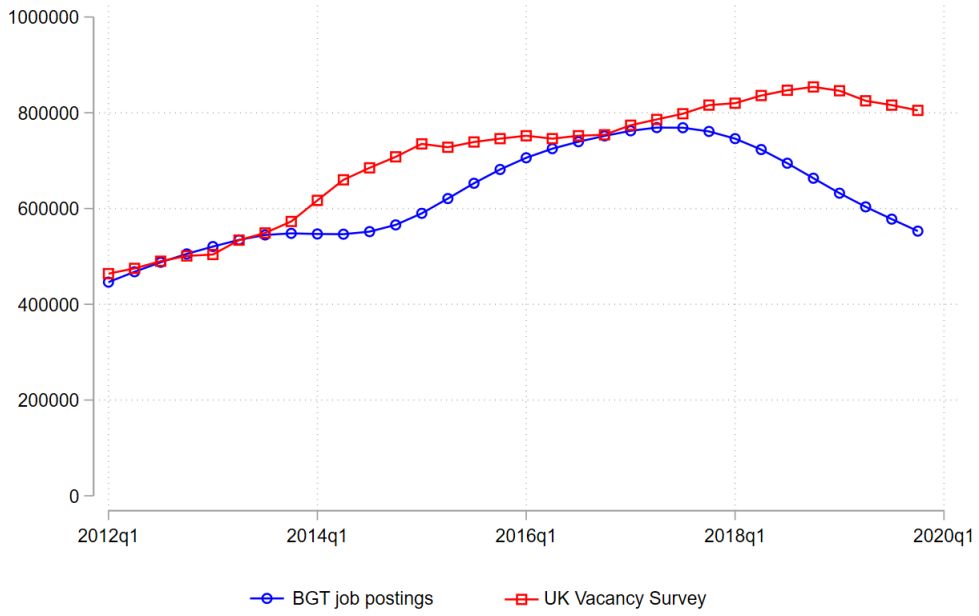


Figure 2.12: Average monthly vacancies: BGT and UK Vacancy Survey

Notes: This figure compares the seasonally adjusted average monthly number of job postings from BGT with the UK Vacancy Survey. We deseasonalise the BGT data using the Hodrick–Prescott high-pass filter.

The total sample is approximately 6,000 businesses per month, with approximately 1,300 large businesses included every month and the remaining 4,700 consisting of smaller enterprises randomly sampled on a quarterly basis. The ONS then constructs total estimates for the UK by weighting the data using employment estimates. For non-responding firms a link factor is calculated and applied to previous returns. The original construction for a never-responding business is calculated from a ratio (calculated from other respondent values in the same sampling strata) being applied to the register employment. For subsequent periods, imputed values are then based on movements in similar-sized businesses. They then provide quarterly seasonally adjusted estimates of the monthly average number of vacancies for the UK economy.[42]

To compare figures from the two sources, we deseasonalise the BGT data and take quarterly averages of the total sum of all postings. Figure 2.12 compares the time series of BGT with the UKVS over the period of 2012-2019. Over this period,

[42]Results are seasonally adjusted in X-12 ARIMA using a multiplicative model.

BGT covers on average 86% of the total reported in the VS. The monthly average number of postings in the BGT data is very similar to the monthly average number of vacancies reported in the VS during two period: 2012 & 2013 and 2016, but the two series diverge during 2014, 2015 and after 2016. This divergence could reflect the fact that the UKVS relies on imputed values, or that firms inaccurately report their number of openings in specific time periods, for example due to uncertainty, or that there has been a structural change in the average number of jobs advertised in one online posting during these time periods. Methods of imputation may not be responsive to labour market trends in the way that directly observed, scraped data would be. We hence might expect the UKVS to over-estimate vacancies in periods of uncertainty on in downturns; this may explain the divergence after the Brexit referendum.

2. ONS Workforce Jobs (WFJ)

We also consider the relationship between labour demand as captured by job adverts and the total number of jobs in the UK economy as measured by the ONS WFJ, which is also affected by labour supply conditions. Vacancies are converted to hires provided the vacancies can be filled. Hires in turn contribute to aggregate job statistics, which additionally depend on the existing number of jobs and firing. The ONS provides a quarterly measure of the estimated total number of jobs in the UK, defined as the sum of employee jobs measured primarily by employer surveys, self-employment jobs from the Labour Force Survey (LFS) and government-supported trainees and Her Majesty's Forces (HMF) from administrative sources.[43] The definition of an employee is anyone working on a specific date who is aged 16 years and over, that is paid in return for carrying out a full-time or part-time job or being on a training scheme. Private sector employee jobs are obtained from Short-Term Employer Surveys (STES), which are conducted with approximately 32,800 businesses per quarter. The Quarterly Public Sector Employment Survey (QPSES) measures public sector employment and is conducted with 1500 contributors per quarter. The LFS surveys 50,000 households per quarter.

[43]https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/methodologies/workforcejobsqmi

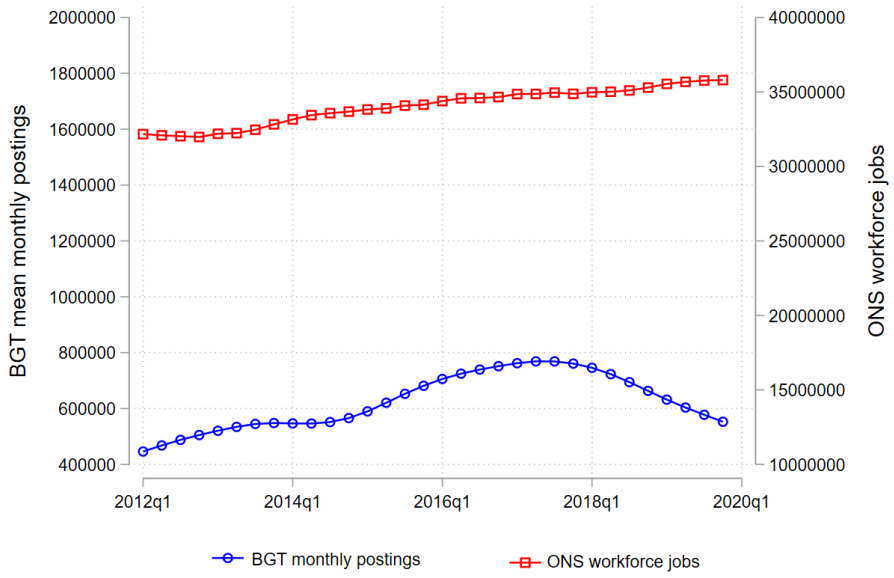


Figure 2.13: BGT average monthly vacancies and ONS average monthly jobs

Notes: This figure compares the seasonally adjusted average monthly number of workforce jobs from the ONS workforce jobs survey. We deseasonalise the BGT data using the Hodrick–Prescott high-pass filter.

Over the period of 2012-2019 the median number of workforce jobs was 34 million. The average ratio of monthly vacancies from the UKVS to the total number of workforce jobs was 2%, while the ratio of job postings from BGT to the total number of workforce jobs was 1.8%. Figure 2.13 displays the total number of jobs over time, compared to the monthly average number of BGT postings over time for each quarter. The number of workforce jobs has steadily increased from 2012 to 2019 and experienced no dip after the referendum.

There are a number of reasons why the WFJ estimates might not have demonstrated a decline after the referendum. If firms do not hire, but also do not fire, then we would not expect the total number of jobs to decrease. This could have been the case after the referendum: due to the uncertainty about future policy, firms and workers may have been unwilling to hire and fire, or leave and search for new jobs. Another possibility is that these statistics also rely on imputation for non-responsive and non-sampled firms and workers and so could be less responsive in times of downturn or uncertainty.

2.B Additional maps and figures

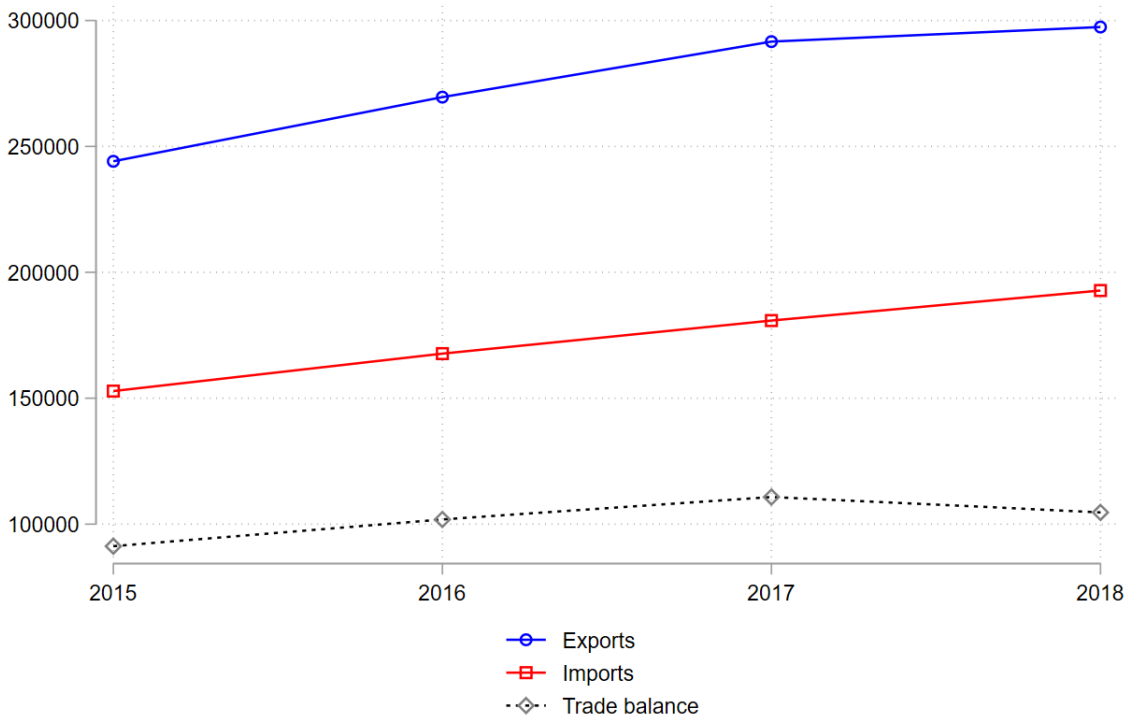


Figure 2.14: UK total trade in services in £millions

Notes: This figure shows the UK's total exports and imports of services, by year, in millions of pound sterling, taken from the UK's Pink Book 2019 (ONS, 2019).

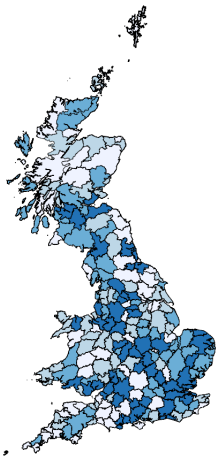


Figure 2.15: Financial service exposure map

Notes: This map displays the measure of exposure to financial services trade barriers by TTWA. Areas with darker blue had a higher employment share in 2015 in sectors more exposure to potential future trade barriers on financial services exports to the EU.

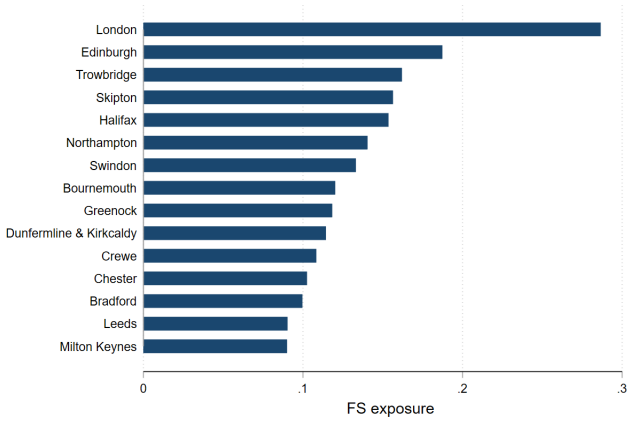
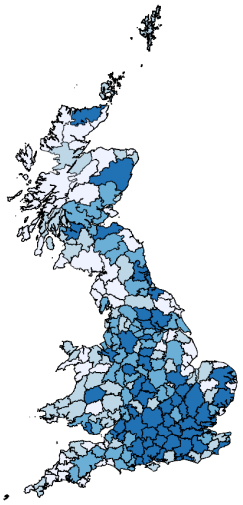
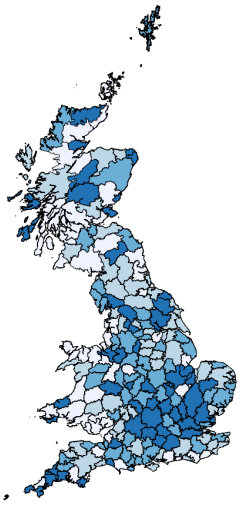


Figure 2.16: Most exposed TTWAs: Financial Services

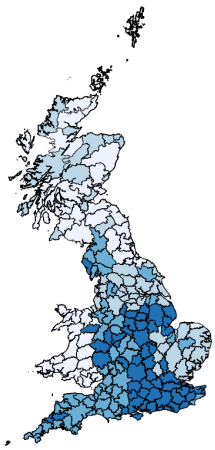
Notes: This chart presents the employment-weighted financial services exposure measures for the top 15 most exposed TTWAs.



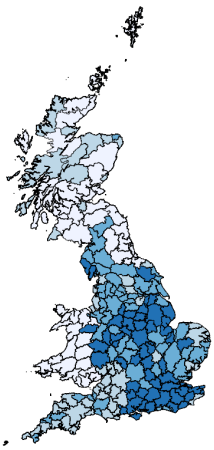
(a) Exposure to more expensive imports



(b) Exposure to cheaper exports



(c) EU national employment share



(d) EU8 national employment share

Figure 2.17: Exposure to the exchange rate depreciation and immigration

Notes: Panel (a) displays the measure of exposure to the impact of the exchange rate depreciation on the price of intermediate imports. Areas with darker blue had a higher employment share in 2015 in sectors more affected by higher cost imports. Panel (b) displays the measure of exposure to the impact of the exchange rate depreciation on the price of exports. Areas with darker blue had a higher employment share in 2015 in sectors with more competitive export prices after the vote. Panel (c) displays the 2015 employment share of EU nationals, areas with darker blue had a higher share. Panel (d) displays the share of EU8 nationals, areas with darker blue had a higher share.

2.C Additional tables

TABLE 2.13. MAPPING BETWEEN UK SIC, OECD STRI & ONS SERVICE EXPORT PRODUCT TYPES

UK SIC 2007 code	STRI sector	ONS service product type
692	Accounting	Accountancy, auditing, bookkeeping and tax consulting services
691	Legal	Legal services
61	Telecom	Telecommunication services
62	Computer	Computer services
63	Computer	Information services
7111	Architecture	Architectural services
7112	Engineering	Engineering services
65	Commercial banking	Financial services
64	Insurance	Insurance services

Notes: This table displays the UK SIC codes with their mapped OECD STRI sectors and ONS product categories used to construct the professional service barrier threat measures.

TABLE 2.14. US PLACEBO FOR FINANCIAL SERVICES

Dep variable: log US postings	(1)	(2)	(3)	(4)	(5)	(6)
	Google		Newspaper		BUI	
measure* FS exposure	-0.010	-0.012	-0.011	-0.010	-0.003	-0.004
	(0.036)	(0.034)	(0.016)	(0.015)	(0.004)	(0.004)
measure *tariff exposure	-0.001	-0.001	-0.000*	-0.000**	-0.000	-0.000
	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
measure*EU national share	-0.032		0.005		-0.010	
	(0.055)		(0.034)		(0.007)	
measure*EU8 national share		-0.054		-0.003		-0.018
		(0.083)		(0.050)		(0.011)
Observations	8,520	8,520	8,520	8,520	8,520	8,520
Adjusted R-squared	0.994	0.994	0.994	0.994	0.994	0.994
TTWA FE	YES	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with various uncertainty measures. The period September 2016 to December 2019 is covered. All specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 2.15. UNCERTAINTY MEASURES WITH QUARTERLY TREND CONTROLS

Dep varable: log postings		Google			Newspaper			BUI	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
post vote*measure*	-0.172***	-0.184**	-0.044***	-0.060***	-0.038	-0.026***	-0.013***	-0.041***	-0.007**
prof service exposure	(0.047)	(0.079)	(0.016)	(0.022)	(0.033)	(0.009)	(0.003)	(0.008)	(0.008)
post vote*		0.049			-0.174			1.448***	
prof service exposure		(0.218)			(0.206)			(0.362)	
post vote*measure*	0.009	0.009	0.009	0.006	0.006	0.006	0.001	0.001	0.001
tariff exposure	(0.009)	(0.009)	(0.009)	(0.005)	(0.004)	(0.004)	(0.001)	(0.001)	(0.001)
post vote*	-57.46	-61.02	-61.05	-76.47	-64.74	-64.16	-34.05	-64.42	-64.85
xrate impact via imports	(50.61)	(51.41)	(51.21)	(50.17)	(51.05)	(50.92)	(51.32)	(50.50)	(50.46)
post vote*	6.335	6.409	6.409	7.590*	7.345*	7.333*	6.464	7.162*	7.172*
xrate impact via exports	(4.039)	(4.033)	(4.034)	(3.997)	(3.999)	(3.999)	(4.159)	(4.102)	(4.106)
post vote*measure*	-0.469*	-0.465*	-0.465*	-0.124	-0.131	-0.131	-0.0178	-0.0150	-0.0150
EU national share	(0.253)	(0.252)	(0.253)	(0.124)	(0.124)	(0.124)	(0.019)	(0.019)	(0.019)
Post vote quarters * prof serv exposure	No	No	Yes	No	No	Yes	No	No	Yes
Observations	12,780	12,780	12,780	12,780	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.985	0.984	0.984	0.985	0.984	0.985	0.985
TTWA FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Year-month FE	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with various uncertainty measures. The quarterly interactions represent the dummies for each specific post-referendum quarter, where Q1 is the third quarter of 2016, interacted with the professional service exposure measure. all specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 2.16. PROFESSIONAL SERVICE EXPOSURE IMPACT BY OCCUPATION

Dep var: log SOC postings	1. Managers, Directors and Senior Officials	2. Professional Occupations	3. Associate Professional and Technical Occupations
post vote * prof services exposure	-0.398***	-0.502***	-0.309*
	(0.148)	(0.140)	(0.162)
post vote * tariff exposure	-0.0222	0.0542	0.00262
	(0.0311)	(0.0456)	(0.0383)
post vote * xrate impact via imports	-66.52	-137.8***	-87.23
	(54.16)	(46.64)	(56.03)
post vote * xrate impact via exports	0.883	12.45***	10.56*
	(4.343)	(4.105)	(5.603)
post vote * EU national share	0.792	0.217	1.039
	(0.863)	(0.805)	(0.963)
	4. Administrative and Secretarial Occupations	5. Skilled Trades Occupations	6. Caring, Leisure and Other Service Occupations
post vote * prof services exposure	-0.0420	-0.286*	0.0261
	(0.152)	(0.165)	(0.128)
post vote * tariff exposure	0.00805	0.00538	0.0346
	(0.0344)	(0.0242)	(0.0336)
post vote * xrate impact via imports	51.05	58.14	-73.94
	(61.15)	(67.99)	(54.85)
post vote * xrate impact via exports	9.613*	1.782	19.07***
	(4.915)	(4.736)	(4.291)
post vote * EU national share	1.213	2.229**	0.323
	(0.962)	(0.989)	(0.818)
	7. Sales and Customer Service Occupations	8. Process, Plant and Machine Operatives	9. Elementary Occupations
post vote * prof services exposure	-0.144	-0.266	-0.254
	(0.145)	(0.160)	(0.159)
post vote * tariff exposure	-0.0361	-0.0390	0.0327
	(0.0244)	(0.0365)	(0.0577)
post vote * xrate impact via imports	-42.19	71.86	46.57
	(54.39)	(65.94)	(62.75)
post vote * xrate impact via exports	2.035	4.032	5.316
	(4.440)	(6.359)	(5.353)
post vote * EU national share	0.996	2.197**	0.889
	(0.921)	(1.025)	(0.926)
Observations	12,780	12,780	12,780
TTWA FE	YES	YES	YES
Month-Year FE	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings by occupation in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable. All specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 2.17. BASELINE EXCLUDING LONDON

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
post vote * prof service exposure	-0.539***	-0.541***	-0.355***	-0.340**	-0.346**
	(0.135)	(0.135)	(0.125)	(0.130)	(0.132)
post vote * tariff exposure		-0.008	0.001	0.003	0.001
		(0.033)	(0.035)	(0.035)	(0.035)
post vote * xrate impact via imports			-74.38	-81.50	-76.98
			(52.61)	(49.95)	(50.52)
post vote * xrate impact via exports			7.904*	9.181**	8.574*
			(4.191)	(4.114)	(4.296)
post vote * EU national share				0.703	
				(0.953)	
post vote * EU8 national share					0.540
					(1.399)
Observations	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.984	0.984	0.984
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable, with and without controls. all specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 2.18. INCLUDING SHARES

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
post vote * prof service exposure	-4.287**	-4.216**	-3.285*	-3.432*	-3.299*
	(1.972)	(1.974)	(1.764)	(1.755)	(1.750)
post vote * prof services emp & export share	0.544*	0.534*	0.419	0.441*	0.422*
	(0.284)	(0.285)	(0.251)	(0.250)	(0.249)
post vote * tariff exposure		-0.032	-0.021	-0.020	-0.021
		(0.060)	(0.060)	(0.060)	(0.061)
post vote * manu emp & exp sh		0.001	0.001	0.001	0.001
		(0.002)	(0.002)	(0.002)	(0.002)
post vote * xrate impact via imports			-50.49	-57.74	-53.42
			(54.59)	(52.82)	(52.97)
post vote * xrate impact via exports			7.284*	8.727**	8.051*
			(4.193)	(4.081)	(4.300)
post vote * EU national share				0.788	
				(0.788)	
post vote * EU8 national share					0.614
					(1.331)
Observations	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.984	0.984	0.984
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable, with and without controls. all specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 2.19. ALTERNATIVE TARIFF MEASURES

	Output weighted		Export weighted		Logged tariffs		Regional export weighted	
Dep variable: log postings	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
post vote * prof service exposure		-0.338***		-0.345***		-0.337***		-0.325**
		(0.124)		(0.128)		(0.124)		(0.124)
post vote * tariff exposure	9.830	12.66	4.416	0.498	10.34	13.47	0.000	0.000*
	(18.09)	(19.21)	(7.898)	(9.710)	(19.11)	(20.18)	(0.000)	(0.000)
post vote * xrate impact via imports		-81.68		-80.78		-81.75		-82.25
		(49.71)		(53.22)		(49.68)		(49.56)
post vote * xrate impact via exports		9.458**		9.067**		9.462**		10.04**
		(4.024)		(4.232)		(4.024)		(3.893)
post vote * EU national share		0.655		0.630		0.655		0.667
		(0.797)		(0.802)		(0.797)		(0.791)
Observations	12,780	12,780	12,780	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.984	0.984	0.984	0.984	0.984	0.984
TTWA FE	YES	YES	YES	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on various average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable, with and without controls. all specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 2.20. GOODS NON-TARIFF BARRIERS

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
post vote * prof service exposure	-0.538***	-0.531***	-0.339***	-0.336**	-0.335**
	(0.132)	(0.139)	(0.126)	(0.126)	(0.129)
post vote * tariff exposure		-0.017	-0.013	-0.010	-0.012
		(0.036)	(0.036)	(0.035)	(0.036)
post vote * goods NTBs		0.069	0.117	0.104	0.108
		(0.166)	(0.175)	(0.172)	(0.173)
post vote * xrate impact via imports			-71.45	-77.93	-73.39
			(53.59)	(51.49)	(51.71)
post vote * xrate impact via exports			8.467*	9.477**	8.862**
			(4.269)	(4.178)	(4.324)
post vote * EU national share				0.573	
				(0.770)	
post vote * EU8 national share					0.348
					(1.303)
Observations	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.984	0.984	0.984
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level, including non-tariff barriers, interacted with the post vote dummy variable, with and without controls. all specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 2.21. SPECIFIC TARIFFS AND QUOTAS

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
post vote * prof service exposure	-0.346***	-0.315**	-0.368***	-0.332**	-0.357***
	(0.124)	(0.129)	(0.125)	(0.129)	(0.124)
post vote * ad val tariff exposure	0.003	-0.027	-0.000	-0.037	0.001
	(0.035)	(0.044)	(0.035)	(0.046)	(0.035)
post vote * quota exposure		0.112		0.134	
		(0.140)		(0.155)	
post vote * specific tariff exposure			-1.838	-1.998	
			(1.316)	(1.289)	
post vote * quota or specific tariff exposure					-3.053
					(1.927)
post vote * xrate impact via imports	-81.092	-73.955	-107.184*	-100.899*	-108.547*
	(50.310)	(49.313)	(55.575)	(53.366)	(54.881)
post vote * xrate impact via exports	9.084**	9.690**	7.931*	8.557**	7.695*
	(4.065)	(4.086)	(4.283)	(4.230)	(4.264)
post vote * EU national share	0.639	0.637	0.837	0.851	0.733
	(0.797)	(0.798)	(0.777)	(0.777)	(0.791)
Observations	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.984	0.984	0.984
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level, including barriers due to specific tariffs and quotas, interacted with the post vote dummy variable, with and without controls. all specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 2.22. IMPORT PROTECTION AND IMPORTED INPUTS

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
post vote * prof service exposure				-0.337***	-0.333**
				(0.125)	(0.129)
post vote * intermediate import tariff exposure	-0.024		0.082	0.105	0.093
	(0.124)		(0.216)	(0.206)	(0.201)
post vote * import protection tariff exposure		-0.008	-0.054	-0.047	-0.045
		(0.017)	(0.035)	(0.034)	(0.034)
post vote * export tariff exposure			0.072	0.063	0.059
			(0.062)	(0.062)	(0.062)
post vote * xrate impact via imports				-82.73	-78.34
				(49.81)	(50.21)
post vote * xrate impact via exports				9.183**	8.693**
				(4.112)	(4.300)
post vote * EU national share				0.726	
				(0.790)	
post vote * EU8 national share					0.701
					(1.285)
Observations	12,780	12,780	12,780	12,780	12,780
Adjusted R-squared	0.984	0.984	0.984	0.984	0.984
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on various average trade barrier exposure measures at the TTWA level interacted with the post vote dummy variable, with and without controls. all specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 2.23. PRE-VOTE PLACEBO FOR 2015

Dep variable: log postings	(1)	(2)	(3)	(4)	(5)
post June '15 * prof service exposure	-0.037	-0.027	-0.068	-0.081	-0.082
	(0.176)	(0.172)	(0.150)	(0.151)	(0.156)
post June '15 * tariff exposure		0.044	0.065	0.058	0.064
		(0.040)	(0.038)	(0.035)	(0.038)
post June '15 * import appreciation			58.58	76.93	63.54
			(63.95)	(60.74)	(59.90)
post June '15 * export depreciation			14.03**	10.86*	12.79**
			(5.734)	(5.033)	(5.264)
post June '15 * EU national share				-1.696	
				(1.152)	
post June '15 * EU8 national share					-0.988
					(1.999)
Observations	2,556	2,556	2,556	2,556	2,556
Adjusted R-squared	0.985	0.985	0.986	0.986	0.986
TTWA FE	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES

Notes: This table displays the results from the regressions of the log of monthly job postings in each TTWA on the average trade barrier exposure measures at the TTWA level interacted with the post dummy variable, with and without controls. The period considered is 2015 and the post dummy takes value one for months July-December. all specifications include TTWA and month-year fixed effects, and standard errors (in parentheses) are two-way clustered at the TTWA and month-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

CHAPTER 3

Artificial Intelligence and Services Offshoring

Abstract This paper constructs detailed measures of the demand for Machine Learning (ML) skills in UK firms, local labour markets and industries using a dataset of the majority of all job adverts posted online in the UK between 2012 and 2017. The demand for ML skills has risen rapidly in this period and the growth has been prevalent across a wide range of industries. I combine these measures with data on UK services trade to explore how the deployment of ML in the UK has affected services offshoring to developing countries. By exploiting technical advances in the task-specific capabilities of ML, which have affected some occupations relatively more than others, I show that ML deployment has led to an increase in services offshoring, particularly to developing countries, rather than high-income countries. This increase has mainly been in ICT-related and Business and Professional Services and India has been the main beneficiary country. In the UK, ML deployment has had a weakly positive impact on labour demand for non-ML roles, particularly for lower skilled occupations.

3.1 Introduction

Advances in Artificial Intelligence (AI) have profound and diverse implications. There has been growing public debate and a rapidly expanding academic literature exploring the consequences of AI for labour markets, productivity and economic growth, particularly in high income countries. Less frequently discussed is the impact of AI on international trade, the geography of global supply chains or the future pathways to economic growth for developing countries.

The potential labour replacing effects of AI, or ML, have received a lot of attention. For example, Frey and Osborne (2017), Brynjolfsson et al. (2018), Felten et al. (2019) and Webb (2019), along with an expanding number of consultancy reports, have attempted to evaluate which tasks or occupations will be replaced by AI in coming years. There has been growing speculation that advances in AI threaten to lead to the reshoring of a great number of jobs in countries like India and the Philippines that have specialised in export-oriented services, such as Business Process Outsourcing (BPO), that often involve tasks that are increasingly feasible to automate using certain applications of ML, such as Optical Character Recognition (OCR) or Natural Language Processing (NLP). For example, Paul Krugman received a lot of attention for warning Indians at the Rising India Summit in 2018 that, 'There is this concept called artificial intelligence that you should be wary of. In future, while diagnosis may be outsourced to a doctor in India, it could also go to a firm based on artificial intelligence' (The Financial Express, 2018).

Reshoring of these activities could have substantial ramifications. In light of the declining labour share of income in manufacturing, many have been looking to the potential of an alternative services-led development model, such as that followed by India. The rise in service sector offshoring from advanced economies over the last 20 years has been a substantial force for employment growth and poverty reduction. In India, for example, the BPO sector currently employs around 4 million people and contributes 8% of GDP and in the Philippines it employs around 1.5 million people and constitutes 17% of GDP (World Bank, 2018; NASSCOM, 2018). Additionally, labour market opportunities in the BPO sector have been shown to be associated with improved schooling, work, marriage and health outcomes for

young women (Jensen, 2010, 2012) and higher primary school enrolment rates (Oster and Steinberg, 2013).

In contrast to the focus on labour replacement, however, an alternative view is that AI will play an enabling role in reducing trade costs and will increase outsourcing. Brynjolfsson et al. (2019) have shown that machine translation can increase trade between individuals who speak different languages. Baldwin (2019) and Baldwin and Forslid (2020) have suggested that AI and digital platforms will facilitate outsourcing of tasks and jobs, making services more easily tradable and that developing countries are likely to benefit. There is also a broader literature documenting how advances in communication technologies previously facilitated the global fragmentation of trade and enabled the most recent wave of globalisation (e.g. Steinwender (2018), Fort (2017) and Baldwin (2016)).

Additionally, there is growing evidence that AI may in fact be quite different to other automation technologies and exhibits the characteristics of a General Purpose Technology (GPT), one that is widely applicable across industries and that catalyses subsequent innovation or is the 'invention of a method of invention' (Cockburn et al., 2018; Goldfarb et al., 2020). In this respect, AI could generate new inventions, new demand and and new tasks, all of which could have the potential to increase offshoring to developing countries.

On the surface, it doesn't appear that the labour replacing effect of AI is the dominant force at play when it comes to services offshoring. In the UK, for example, service imports from India, which is the UK's largest services offshoring destination, more than doubled over the past decade. Revenues in India's BPO sector nearly tripled over the past ten years (NASSCOM, 2018). Large offshore outsourcing companies such as Tata Consultancy Services or Infosys, that initially provided BPO services focused on labour intensive, repetitive tasks, such as invoice processing or data entry, have continued to grow and now provide a whole range of advanced IT outsourcing activities.

The adoption of AI in India has also been increasing rapidly. For example, India was one of the five countries with the fastest growth in AI hiring between 2015 and 2019 according to the Stanford AI Index (2019). There are also plenty of examples of European or US firms offshoring the task of developing AI-driven

automation projects to IT-outsourcing companies in India. For example, Infogain, a Delhi-based Indian IT company, provides specialist AI and robotic process automation (RPA) services to a host of US based multinationals such as Airbnb and Apple, assisting these companies to automate business functions conducted by their workers in the US.

Beyond these aggregate statistics or anecdotal observations, however, empirical evidence on this topic is scarce. Much of the research on AI to date has focused on the supply side, for example studying academic research in the field of AI, patenting or startups (e.g. Klinger et al. (2018), Stanford AI Index (2019)). There is far less data or research on the commercial uptake of AI. In fact, we know extremely little about the extent to which AI is being deployed commercially, or for what purposes. As such, there have been some claims that advances in AI have been exaggerated and there has been little real world application. For example, the State of AI Report by MMC Ventures and Barclays (2019) captured a lot of attention by suggesting that of 2830 AI startups surveyed in the EU, around 40% were not actually using AI. Detailed empirical research with measurement of the commercial deployment of ML is required to evaluate diffusion of AI adoption and investigate its complex different impacts.

In this paper I contribute to this knowledge gap by developing new proxy measures of ML deployment at the firm, local labour market and industry level in the UK. I take inspiration from Tambe and Hitt (2012), who demonstrate that technology diffusion can be measured using labor demand data, and construct measures of the demand for ML skills and expertise as listed in online job adverts. I do so using a comprehensive scraped dataset of nearly all of the job adverts posted online in the UK between 2012 and 2017. This allows me to construct detailed measures of the demand for specific skills that reflect the deployment of ML.

I document a substantial rise in the demand for ML skills in the UK between 2012 and 2017, casting doubt on the assertions that advances in AI have been exaggerated and are inconsequential in practice. This increase in demand has been prevalent across a wide range of industries, providing suggestive evidence in support of the view that AI is a GPT. The UK local labour markets that have experienced the greatest rise in demand have been Cambridge, Oxford, London, Read-

ing and Bristol, while the 2 digit Standard Industrial Classification (SIC) industries with the greatest demand for ML skills have been financial services, retail trade, legal and accounting, management consultancy and office administrative and support activities.

I combine these measures of the demand for ML skills with a detailed dataset on services trade collected from the UK's Inquiry on International Trade in Services (ITIS), which allows me to construct measures of services offshoring by country and service type. The ITIS dataset lends itself well to studying services offshoring because it measures imports of a number of categories of services that are used as intermediate service inputs, such as legal services, accounting services or telecommunication services. I use the metrics constructed by Felten et al. (2019), which measure the extent to which different occupations are exposed to automation using ML, to construct measures of *ex ante* exposure of firms, local labour markets and industries to advances in ML. Felten et al. (2019)'s measures were developed by combining data from the Electronic Frontier Foundation (EFF)'s AI Progress Metrics, which track progress in specific AI capabilities, such as image recognition or machine translation, with data on the task content of occupations as documented by O*NET.

I show that firms, local labour markets and industries that were initially more exposed to being automated using ML subsequently experienced a greater increase in the posting of job adverts demanding ML skills between 2012 and 2017. This suggests that opportunities to automate tasks with ML are being exploited by firms through investment in the human capital required to build or deploy ML systems. Using baseline ML exposure as an instrument for the growth in ML demand, I then explore the impact of ML demand on other types of labour demand in the UK and on services offshoring.

I find that within UK firms, growth in demand for ML skills has had a positive impact on the growth in demand for lower skilled occupations, but has had no effect on the growth in demand for higher skilled occupations. There is some evidence that this former effect dominates and the net impact on total non-ML vacancies within firms is positive, particularly when comparing firms within the same industry. Growth in demand for ML skills has had a negative effect on the

demand for the occupations that are the most potentially replaceable by ML, suggesting that substitution does occur between occupations that can be automated using ML and ML expertise. However, there is some evidence that the growth in demand for ML skills has had a larger, positive impact on the posting of vacancies for non-ML exposed occupations.

I next look at the impact of ML demand on services offshoring at the local labour market and industry level. I find that growth in ML demand actually caused an increase in services offshoring, with particular evidence of an increase in offshoring to low and middle income countries, rather than high income countries. The most positively affected country has been India, with some evidence of a positive effect for South Africa, the UK's fourth largest services offshoring destination. The increase in offshoring has been concentrated in the offshoring of business services and ICT-related services, such as computer, information and telecoms services. These findings suggest that, while there is evidence of a displacement effect of AI for the most exposed occupations within the UK, the productivity or enabling effects of AI for services offshoring have dominated, particularly in terms of offshoring to developing countries.

This paper makes a number of new contributions to the emerging literature studying AI, automation and trade. First, I add new evidence to the small but growing literature looking directly at how automation in high income countries affects trade with developing countries (Artuc et al., 2019, 2018; Pedemonte et al., 2019; Faber, 2018; Hallward-Driemeier and Nayyar, 2019). These papers have all studied robots and goods trade; I add a new angle by studying AI and services trade. My results in this different context of AI and services offshoring are generally consistent with the findings that automation in high income countries, in fact, increases imports from developing countries.

Second, I contribute to the literature studying how AI and other digital or communication technologies enable trade and the fragmentation of global supply chains (Brynjolfsson et al., 2019; Freund et al., 2019; Goldfarb and Trefler, 2018; Fort, 2017; Baldwin and Forslid, 2020). While I cannot disentangle the mechanisms driving the results in this paper, it offers some initial empirical evidence that AI

could be an enabling force for services offshoring, particularly to developing countries.

This paper also contributes to the expanding literature studying the economic impacts of AI (Agrawal et al., 2019; Aghion et al., 2017; Rock, 2019; Klinger et al., 2018). In particular, I add novel empirical evidence to the literature studying the impacts on labour demand. While there have been many papers studying which tasks and occupations are technically feasible to automate with AI, there has been scant evidence on what is actually being automated in practice. I demonstrate that one of the leading measures of the replaceability of occupations with AI, that of Felten et al. (2019) is indeed a good predictor of commercial application of AI, as measured using the demand for ML skills.

I also shed light on how ML deployment is affecting other forms of labour demand. The results in this paper provide some evidence that AI is being used to substitute for replaceable occupations, consistent with the predictions of the canonical task-based framework of Acemoglu and Restrepo (2018) and with the suggestions in Acemoglu and Restrepo (2019b) that many of the current applications of ML are focused on automation. However, this paper also shows that ML deployment has increased demand for lower skilled and non-replaceable occupations, with some evidence of a net positive impact on total labour demand, suggesting that some labour reinstating effects of AI have also been at play. In addition, I also contribute to the literature studying whether AI is a GPT by documenting the prevalence of ML skills demand across a range of different industries.

Finally, this paper also adds new evidence to a growing literature studying service sector trade (Breinlich and Criscuolo, 2011; Ariu et al., 2020, 2019; Eppinger, 2019; Liu and Trefler, 2019). I shed some new light on the role of AI in affecting services offshoring, demonstrating that it can play an important role in influencing services supply chains.

This rest of this paper proceeds as follows. Section 3.2 provides background information on services trade, AI and their relationship. Section 3.3 describes the data used, Section 3.4 outlines the measurement of ML skills, Section 3.5 provides the empirical specification, Section 3.6 provides the results, Section 3.7 provides robustness checks and Section 3.8 concludes.

3.2 Background

3.2.1 Services trade and offshoring

Services trade is less widely researched and less easily measured than goods trade. In this paper I define trade in services in accordance with the residential definition of the IMF Balance of Payments Manual, which also underlies the compilation of balance of payments statistics in the UK (see IMF (1993)). International trade in services is defined as service transactions between residents and non-residents of an economy. This definition includes three of the four modes described in the General Agreement on Trade in Services: cross-border supply (mode 1), consumption abroad (mode 2), and presence of natural persons (mode 4). In this definition the concept of residence is key. While the subsidiary of a US multinational in the UK might be foreign-owned, it is ordinarily resident in the UK. As such, its transactions with other UK firms or local consumers do not count as service trade under this definition.

ITIS, the main data source used to measure offshoring, described in detail below, does not cover the entirety of the UK's international services transactions. It focuses primarily on producer, or intermediate, services. It excludes consumer services such as travel, passenger transport and higher education. The method of collection for data on services trade differs substantially from those of goods trade data; ITIS is entirely survey based.

There are a range of definitions of offshoring that have been used in the trade literature. Blinder and Krueger (2013) define offshoring as 'The movement of home-country jobs to another country'. Hummels et al. (2014) define it as 'The import of intermediate inputs that could have been produced within the same firm'. When studying the offshoring of services, I am therefore interested in imports of intermediate services that could have been produced within the same firm. ITIS, which focuses on producer and intermediate services, hence lends itself well to studying offshoring. I focus on a subset of categories of service imports, for example the processing of insurance claims, that are used as inputs into the firm's wider activities conducted in the UK. I discuss this in more detail in Section 3.3.

3.2.2 AI

AI is defined by the Oxford English Dictionary (2020) as 'The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages'. The scope of AI is broad and while some researchers are working on 'Artificial General Intelligence', much of the recent progress in AI has been in the 'narrow AI' subfield of ML. There is no widely accepted definition of ML among practitioners, but one definition is 'the statistical techniques that enable computers and algorithms to learn, predict and perform tasks from large amounts of data without being explicitly programmed' (Acemoglu and Restrepo, 2019b).

ML algorithms are typically divided into two categories: supervised and unsupervised algorithms. Supervised ML occurs when there is a training dataset with input variables and an output variable and an algorithm is trained to learn the mapping function from the input to the output. Once the mapping function has been learnt it can then be applied to new input data in order to 'predict' the output variables for that data. For example, a supervised learning algorithm may learn a function that maps the input of an image to the output of a textual description of the image, using a training dataset of examples of images and their descriptions. Unsupervised ML, on the other hand, occurs where there is only the input data and no corresponding output variables.

Advances in ML have therefore recently been conceptualised as reducing the cost of the task of prediction, which is prevalent in many occupations (e.g. Agrawal et al. (2018)). For example, a radiologist takes the input of medical images to generate the predicted output of a medical diagnosis or an offshored back office employee for a multinational bank takes the input of scrawled handwriting on a mortgage application form to generate a predicted output of the correctly spelt name of the applicant. In this regard, advances in ML could be seen as providing lower-cost or higher-quality inputs into many existing production processes, in a similar way to other automation technologies, like robots. Some early work has hence started to model ML in a similar way to other forms of automation in

the canonical task-based framework of Acemoglu and Restrepo (2018) (e.g. Webb (2019)).

However, in light of the fact that prediction tasks are pervasive in business and society, it has also been suggested that AI may be a GPT (e.g. Brynjolfsson et al. (2017); Goldfarb et al. (2020); Cockburn et al. (2018)). The concept of GPTs originates from Bresnahan and Trajtenberg (1995, p.84), who define them as 'characterized by the potential for pervasive use in a wide range of sectors and by their technological dynamism'. Bresnahan (2010, p.764) further elaborates that a GPT '(1) is widely used, (2) is capable of ongoing technical improvement, and (3) enables innovation in application sectors'. Early empirical evidence from Klinger et al. (2018) and Goldfarb et al. (2020) suggests that AI does indeed exhibit these characteristics, is being deployed across widely different industries and shows signs of being the 'invention of a method of invention'. AI might therefore be expected to have a distinct economic impact to other labour-replacing technologies.

3.2.3 AI, labour demand and services offshoring

Concerns about the impact of AI on labour markets or offshoring typically stem from focus on the direct task displacement effects of AI. In the context of services offshoring, for example, back-office activities that have typically been offshored to lower wage countries, such as processing insurance claims, managing payroll and data entry, tend to be intensive in tasks that are simple prediction problems that could be automated using supervised ML. There are a number of other ways in which AI could effect labour demand and offshoring, however, as hinted at above. The relationship is likely to be complex. Based upon the early conceptual frameworks and research that exist to date, I outline some of the key potential impacts:

1. **Displacement effect**: the direct replacement of tasks previously conducted by humans now conducted by algorithms, as in Acemoglu and Restrepo (2018) or Chapter 1 of this thesis if also incorporating firm heterogeneity and offshoring.

2. **Productivity effect**: the effect of reducing the costs or increasing the quality of factor inputs, which could increase the demand for labour, as in Acemoglu

and Restrepo (2018) or Chapter 1 of this thesis if also incorporating firm heterogeneity and offshoring.

3. **Innovation effect**: an economy wide effect of the improved ability to discover patterns and conduct research. Another way to view this within the task based framework would be if there were previously tasks that were so prohibitively costly for humans to perform that they were never conducted. Now with algorithms they can be conducted, allowing for new demand, new products or services. This could also be called the 'GPT effect' or 'enabling effect' of AI for labour as described in Acemoglu and Restrepo (2019b).

4. **Enabling effect** for trade: Reduced trade costs through applications of ML like machine translation, optimised supply chains or improved matching between buyers and suppliers.

In the context of services offshoring, AI might additionally lead to increases in offshoring if firms seek to offshore parts of the process of developing or deploying ML systems, for example to an IT outsourcing company in India. The net effect of AI on labour demand and services offshoring may therefore depend on the relative importance of these factors. The empirical analysis in this paper will attempt to shed some light on which of these effects dominate.

3.3 Data

3.3.1 Job postings data

The primary dataset used is provided by Burning Glass Technologies (BGT), a private market research firm, that scrapes, parses and deduplicates online job postings and creates labour market analytics products using this data. BGT identify that their postings are sourced from approximately 40,000 online job boards and company websites, resulting in a sample of nearly all job adverts posted online. For the UK there are around 50 million UK job adverts over the period 2012-2017. The nature of online job adverts is that the information they include varies substantially. Some are extremely detailed and include name, address, comprehensive requirements and responsibilities, while others might only include one line.

BGT aim to standardise the adverts by using text analysis to classify them into UK SIC codes and Standard Occupational Classification (SOC) codes and by extracting and classifying information included, such as the skills listed, the experience required, qualifications demanded and location.

This dataset hence has the advantage that it includes rich and detailed information on a number of dimensions that are not possible to obtain from other traditional sources of labour market data. Particularly useful for this paper are the skills listed in the job adverts and the SOC codes. In Section 3.4 I will describe in detail how I construct measures of the demand for ML skills using BGT's skill classification. This detailed text data on the exact requirements listed in the job adverts make it possible to measure the demand for narrowly defined categories of technology-related expertise and measure the diffusion of technology. This would not be possible using traditional data, such as labour force surveys, that include dimensions such as SOC codes, that are more broadly defined and only updated relatively infrequently. For example, in the UK's SOC hierarchy, the closest 4 digit categories to ML engineers might be 'IT business analysts, architects and systems designers' or 'Programmers and software development professionals'. Having information that a company employs these occupations would be insufficient to induce that they are deploying ML.

An additional benefit is that the inclusion of SOC codes allows for a detailed picture of the occupational composition of the labour force by firm, detailed geography and industry. In the UK this information is hard to come by from other sources and so researchers are typically precluded from the type of detailed local labour market or worker-firm analysis that is often conducted in countries that collect more comprehensive worker-firm data, such as Denmark or Brazil. In this paper I use the SOC codes to construct measures of the baseline occupational composition of firms, local labour markets and industries, allowing me to adopt a Bartik style identification strategy, and also to explore the impacts of ML on different occupations.

Despite these advantages, working with job adverts data has a number of downsides too. This data only covers the subset of all vacancies advertised online and this subset advertised online may vary by occupation, job type, geography, in-

dustry and time. In Chapter 2 of this thesis we compare the number of job adverts included in the BGT data with the ONS estimates of the total number of vacancies in the UK economy, finding that BGT includes 86% of the ONS estimates. This suggests that the share of jobs advertised online is actually relatively high in the UK. The ONS does not include the number of vacancies by occupation or region, so we cannot analyse these aspects of coverage. However, in the US, where this data does exist, Hershbein and Kahn (2018) provide a detailed analysis of the industry-occupation mix of vacancies in BGT relative to other detailed US data sources and how this mix has changed over time, finding that the BGT postings are disproportionately concentrated in occupations and industries that typically require greater skill, but that the distributions are relatively stable across time and the aggregate and industry trends in the quantity of vacancies track official sources reasonably closely.

Another complication is that job postings measure just one margin by which a firm may adjust its labour inputs and reflect stated, but not necessarily realised, demand. Labour demand will also be affected by separations, wage changes and training. Lazear and Spletzer (2012) have estimated that around two thirds of hiring is replacement hiring, while the other third is due to business growth. At the firm level, replacement may happen relatively infrequently and so job adverts may not provide a representative snapshot of the composition of the labour force in a given time period. Additionally, the hiring that reflects business growth will be biased towards new and quickly growing areas of the economy.

In the context of measuring ML demand, I am focusing on a high-skilled white collar role and so the representation of ML skills and occupations in the BGT data could be expected to be high. I attempt to mitigate some of the other potential issues in working with job adverts data in the following ways. First, I measure ML skills as the ratio of ML job adverts to total job adverts in order to reduce the issue of certain firms, regions or industries advertising a lower proportion of vacancies online. This approach therefore assumes that if a firm, region or industry is underrepresented in the BGT data, its ML demand will be just as underrepresented as its demand for other occupations. Second, I employ a long differenced empirical specification to avoid short-run fluctuations in the the demand for certain new or

fast growing industries, occupations or regions. Finally, when measuring baseline occupational shares by firm, I take the shares over a two year period, so as to be more likely to obtain a representative sample of the occupations a firm employs.

Firm level analysis

For the analysis on the impact of ML on labour demand in the UK, I construct a panel of establishment level job adverts using the BGT data only, constructed using the company names and 4-digit postcode areas. The BGT data includes both company name and latitude and longitude for each job advert. I first keep only the job postings with an employer name included. For the period of 2012-2017 these represent 29% of all postings. I then take the names in lower case and remove suffixes, such as 'public limited company' from the end of the company names.[1] I do so because job postings from the same firm sometimes list their name slightly differently by including or excluding these words. I then extract the postcode associated with the longitude and latitude.[2] These postcodes are often only accurate to an aggregated level and so I extract the first four digits from the postcode. I then group together any job advert with the same derived company name in the same 4 digit postcode area and designate these as 'establishments', with their 'parent company' as the group with the same company name but a different 4 digit postcode area.

Classifying job adverts by skill

The BGT data are classified by 4 digit SOC code. Examples include 'Managers and proprietors in agriculture and horticulture' or 'Metal plate workers, and riveters'. These can be aggregated to nine 1-digit groups as presented in Table 3.1.

The ONS classifies the 2-digit sub-major groups of the SOC 2010 into four skill levels where 'skill level is defined with respect to the duration of training and/or work experience recognised in the field of employment concerned as being normally required in order to perform the activities related to a job in a competent

[1]The full list of suffixes removed is: 'public limited company', 'limited', 'ltd', 'ltd.', 'plc', 'p.l.c.', 'inc', 'group', '(group)', 'corporation', 'internatioal', 'uk', '(uk)', '(u.k.)', '(united kingdom)', 'the', '(gb)', 'holdings', '(holdings)' and 'llp'.

[2]To do so I make use of OpenCage geocoder https://opencagedata.com/

TABLE 3.1. POSTING OCCUPATIONS AND SKILL LEVELS

SOC Code	Group title	Skill grouping
1	Managers, Directors and Senior Officials	3/4
2	Professional Occupations	4
3	Associate Professional and Technical Occupations	3
4	Administrative and Secretarial Occupations	2
5	Skilled Trades Occupations	3
6	Caring, Leisure and Other Service Occupations	2
7	Sales and Customer Service Occupations	2
8	Process, Plant and Machine Operatives	2
9	Elementary Occupations	1

Notes: Where a skill grouping of 4 is the highest skill level and 1 is the lowest as defined by the ONS. My definition of high skill includes levels 3 and 4, and low skill includes 1 and 2. *Source*: ONS Standard Occupational Classification.

and efficient manner'.[3]

I define the top two levels (3 and 4) as 'high skill' and the bottom two (1 and 2) as 'low skill'. Examples of high skill sub-major groups include 'Science, research, engineering and technology professionals' (level 4) and 'Business and public service associate professionals' (level 3). Examples of low skill sub-major groups include 'Administrative occupations' and 'Elementary trades, plant and storage related occupations'.

3.3.2 Service sector trade data

The service sector trade data is from ITIS, collected by the ONS (Office for National Statistics, 2019). ITIS provides data on the value of transactions of UK businesses by country of origin, destination and product category. Businesses are legally required to respond to ITIS and ITIS is used in constructing the UK's Balance of Payments statistics and contributes towards the measure of GDP. ITIS consists of two non-overlapping surveys: a quarterly sample of approximately 2,200 businesses and an annual sample of approximately 15,500 businesses. Both are directly sampled from the Inter-Departmental Business Register (IDBR), a live register of UK businesses. The ITIS survey covers the whole economy, with three exceptions; banking, transport and the public sector, which are collected from other sources.

[3]https://www.ons.gov.uk/methodology/classificationsandstandards/standardoccupationalclassificationsoc/soc2010/soc2010volume1structureanddescriptionsofunitgroups

The sampling methodology consists of three parts. First, 'known traders' are selected from the responses of previous years. Second, filter questions in the Annual Business Survey (ABS) are used to identify traders by asking reporting units to indicate whether they imported services or/and exported services. All positive responses not already in ITIS are added. Finally, stratified random sampling is applied to 'high propensity' industries in the IDBR, based on employment-defined strata with sampling fractions decreasing in direct proportion to employment. As with other surveys conducted by the ONS, the survey is sent out to 'reporting units'. In the vast majority of cases a reporting unit is equivalent to a business or enterprise, but large enterprises may have several reporting units.

This setup results in a *de facto* dataset of the transactions of almost all UK firms trading in these services, excluding those in banking, transport and the public sector. ITIS covers 51 product categories. In this paper I am interested in services offshoring, as defined in Section 3.2 and so want to focus on intermediate producer services, rather than consumer services and service sector offshoring, rather than agricultural or manufacturing services. I therefore focus primarily on products in the groups of Business and Professional Services, Telecommunications, Computer and Information Services and Technical and Scientific Services. The full list of services included is found in Appendix 3.A. Examples are accountancy, legal services, telephone services, insurance and pension services and engineering services. The excluded service products are primarily in the categories of Agricultural and Mining Services, Manufacturing, Maintenance and On-site Processing Services, Intellectual Property, Construction, Merchanting and Other Trade-related Services and Personal, Cultural and Recreational Services.

I further break down imports into those from low and middle-income (LMI) countries and high income (HI) countries, as defined by the World Bank in 2012. Figure 3.1 shows the UK's top four LMI country import partners and how imports have been changing over the past decade. India is by far the largest import partner in this category and service imports from India have more than doubled over the past decade. The next largest import partners are China, South Africa and Brazil, which I will focus on in detail in the empirical analysis. After these four the next largest are Romania, Egypt, the Philippines, Bulgaria and Turkey. Figure 3.2 dis-

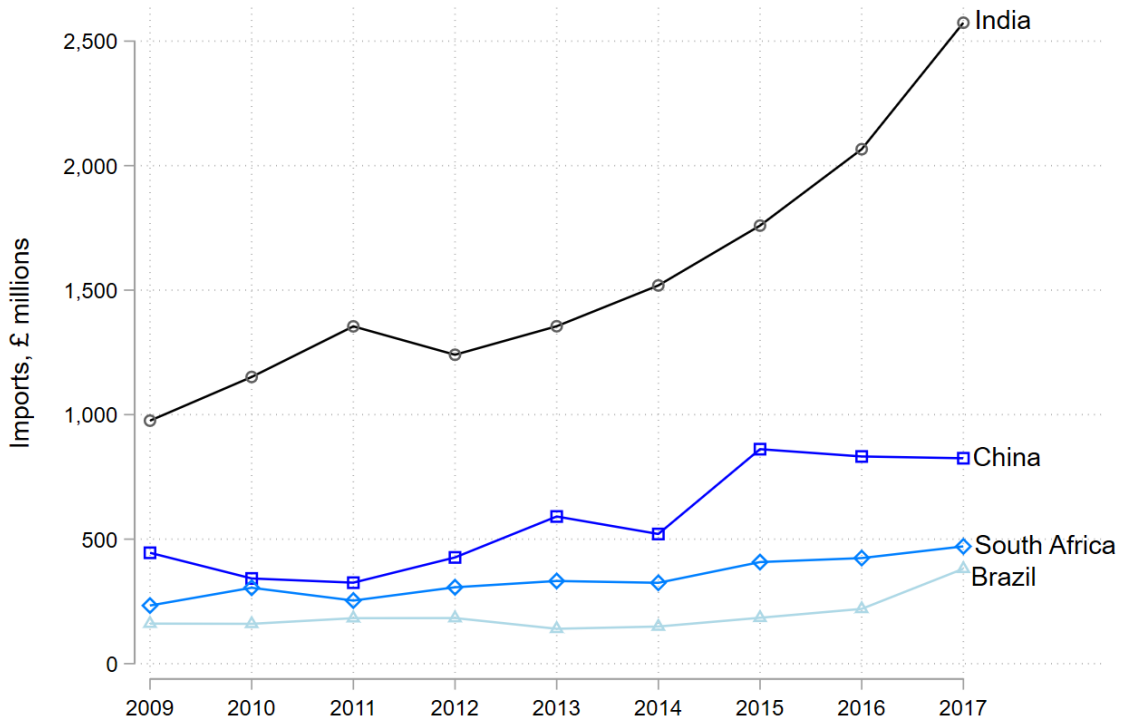


Figure 3.1: UK service imports from top 4 LMI country partners

Notes: This figure displays the total value of UK service imports from the UK's top 4 LMI import partners, over time. *Source:* ITIS dataset.

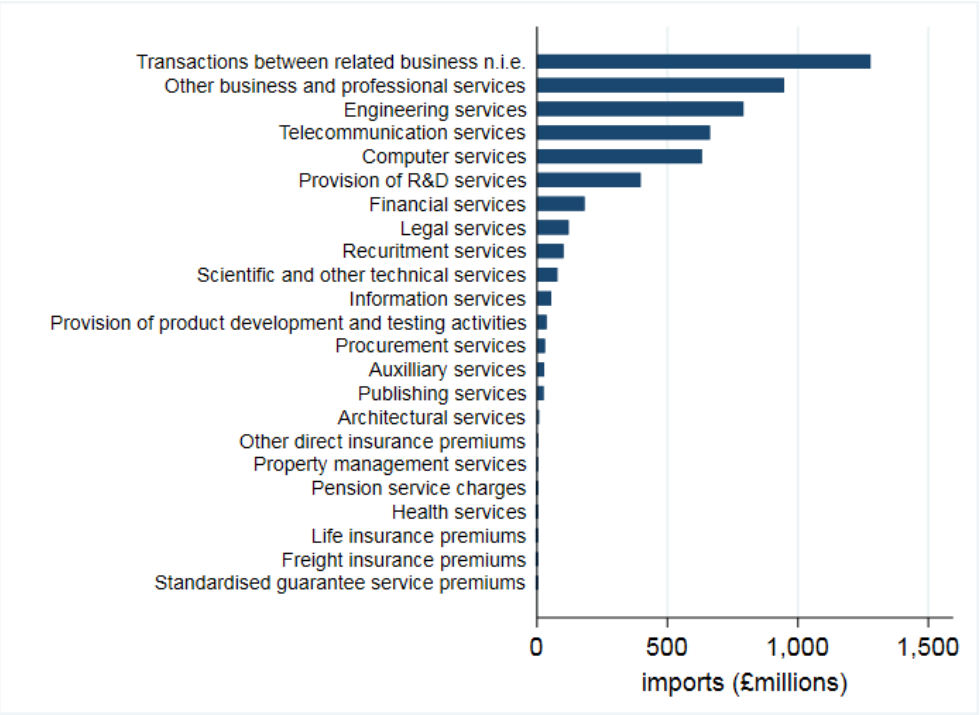


Figure 3.2: UK service imports from LMI countries by product

Notes: This figure shows the total value the UK's service imports included in the ITIS dataset from LMI countries, by service product. *Source:* ITIS dataset.

plays the largest import product categories from LMI countries in 2017. The largest value of imports is in the category of 'Transactions between related businesses not included elsewhere (n.i.e)', followed by 'Other business and professional services', 'Engineering services', 'Telecommunications services' and 'Computer services'.

3.3.3 Additional datasets

I also make use of data from the ABS and the UK Business Register and Employment Survey (BRES). The ABS is a census of all large businesses employing 250 or more, and a sample of smaller businesses (Office for National Statistics, 2020a). BRES provides the number of full and part time employees for each reporting unit at each site location ('local units') (Office for National Statistics, 2020b). I make use of the public/private employment split from BRES and the computer service purchases from the ABS, both used as baseline control variables.

3.3.4 Aggregating data to local labour market and industry level

For the analysis regarding services offshoring, I combine the BGT and ONS datasets by local labour market or industry. For local labour markets I use UK Travel to Work Areas (TTWAs), which aim to reflect the geographic region where the population would generally commute to a larger town, city or conurbation for the purposes of employment.[4] The current criteria for defining TTWAs are that at least 75% of the area's resident workforce work in the area and at least 75% of the people who work in the area also live in the area. The area must also have an economically active population of at least 3,500. TTWAs range in population size from 6,800 to 8.4 million.

I aggregate the ITIS, BRES, ABS and BGT datasets to the 2011 TTWA level and SIC 2007 3-digit level as follows. BRES provides the detailed TTWA of each local unit surveyed, where each local unit is subsidiary to a reporting unit. For each reported unit in ITIS, I distribute the value of services transactions according to the employment shares of local units within that reporting unit. For BRES and the ABS, I make use of the local unit data to construct totals of variables at the TTWA level. BRES and ITIS both include 218 TTWAs. The BGT data also classifies postings by TTWA, with only 5 TTWAs for which there are no postings. The resulting sample consists of 213 matched TTWAs.

For the SIC 3-digit classification, each reporting unit and local unit is already assigned an SIC code and so variables can easily be aggregated to these. The BGT dataset has far more comprehensive coverage of postings by TTWAs than SIC codes, due to the difficulty in assigning an industry to a job advert. For the period of 2012-2017 only 20% of postings are missing a TTWA classification, while 54% are missing an SIC3 code. The TTWA level analysis for offshoring therefore serves as the primary focus, with the industry level analysis a secondary exercise. ITIS includes 215 SIC3 codes, while BRES includes 259. I assign a zero trade in services value to all sectors not included in ITIS but included in BRES. The resulting sample with overlap also with BGT consists of 246 sectors. Table 3.2 displays

[4]TTWAs are defined by the ONS using census data for commuting between wards, based on the different locations of individuals' home and work addresses. See here: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/articles/traveltoworkareaanalysisingreatbritain/2016

TABLE 3.2. SUMMARY STATISTICS

	mean	sd	min	max	obs
	Establishment leve variablesl				
A. Variables in levels (annual)					
Postings	57.4	408.5	1	54416	87298
High skilled postings	37.8	303.1	0	42092	87298
Low skilled postings	19.6	112.6	0	13556	87298
Most exposed postings	34.4	267.5	0	36769	87298
Least exposed postings	6.6	42.7	0	4954	87298
ML postings	0.08	1.2	0	106	87298
Broad ML postings	0.3	2.9	0	219	87298
ML share	0.0009	0.01	0	0.7	87298
Broad ML share	0.003	0.02	0	1	87298
B. Transformed differenced variables (2012&13-2016&17)					
Δ log postings	0.4	1.0	-5.0	4.8	15260
Δ log non-ML postings	0.4	1.0	-4.9	4.8	15260
Δ IHS(high skilled non-ML postings)	0.4	1.1	-4.7	5.9	15260
Δ IHS(low skilled non-ML postings)	0.4	1.2	-7.5	6.3	15260
Δ IHS(most exposed non-ML postings)	0.4	1.2	-5.0	5.7	15260
Δ IHS(least exposed non-ML postings)	0.3	1.4	-8.0	6.5	15260
Δ ML share	0.001	0.01	-0.4	0.4	15260
Δ Broad ML share	0.004	0.03	-0.5	0.7	15260
C. Baseline variables (2012&13)					
Exposure	0.7	0.02	0.6	0.7	15260
Customer service share	0.04	0.1	0	1	15260
Administrative & secretarial share	0.09	0.1	0	1	15260
	TTWA level variables				
A. Variables in levels (annual)					
Postings	27238	118765	22	1935761	426
ML postings	62.3	758.6	0	15418	426
Broad ML postings	198.8	2340	0	47778	426
ML share	0.0004	0.001	0	0.0143	426
Broad ML share	0.002	0.003	0	0.0248	426
Offshoring to LMI countries (£mil)	15.9	97.7	-	-	426
Offshoring to HI countries (£mil)	109.6	657.8	-	-	426
B. Transformed differenced variables (2012-2017)					
Δ ML share	0.0006	0.001	-0.002	0.012	213
Δ Broad ML share	0.002	0.004	-0.01	0.021	213
Δ Log Offshoring to LMI countries	-0.16	.74	-	-	213
Δ Log Offshoring to HI countries	-0.03	0.86	-	-	213
C. Baseline variables (2012)					
Exposure (2012&13)	0.68	0.006	0.66	0.70	213
Customer service share	0.03	0.02	0	0.15	213
Administrative & secretarial share	0.10	0.02	0 .04	0.18	213
Public sector share	0.17	0.08	-	-	213

Notes: This table displays the summary statistics for the key variables used in the firm-level and TTWA level analysis. Potentially disclosive data replaced with '-'. I display only non-ML postings derived for the narrow measure, results for the broad measure are almost identical. Sources: BGT for Establishment level, BGT, ITIS and BRES for TTWA level.

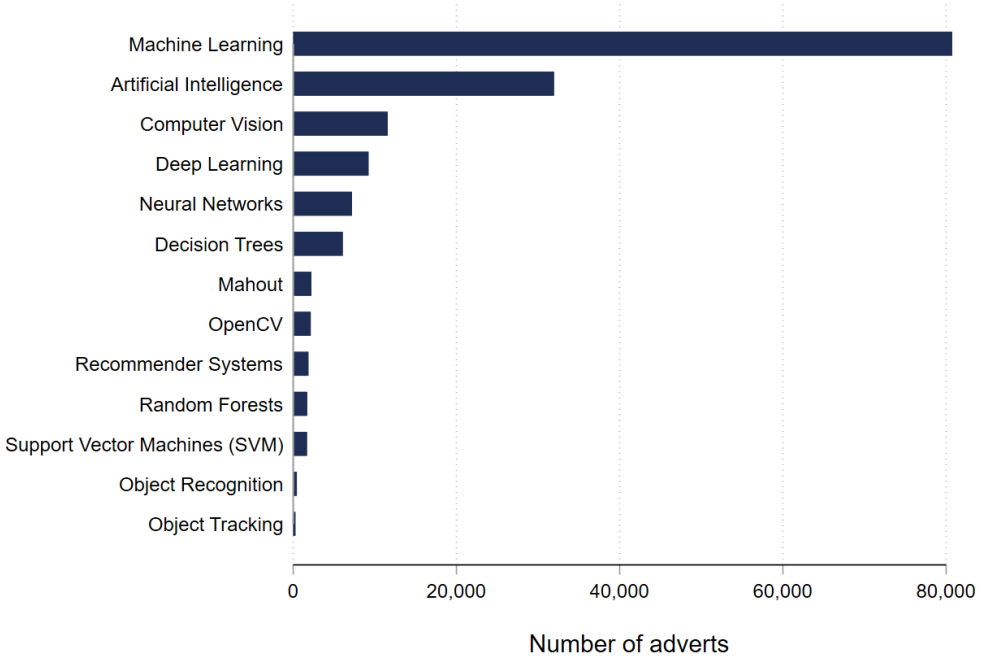
the summary statistics for the key variables used in both the firm level and TTWA level analysis.

3.4 Measuring demand for ML skills

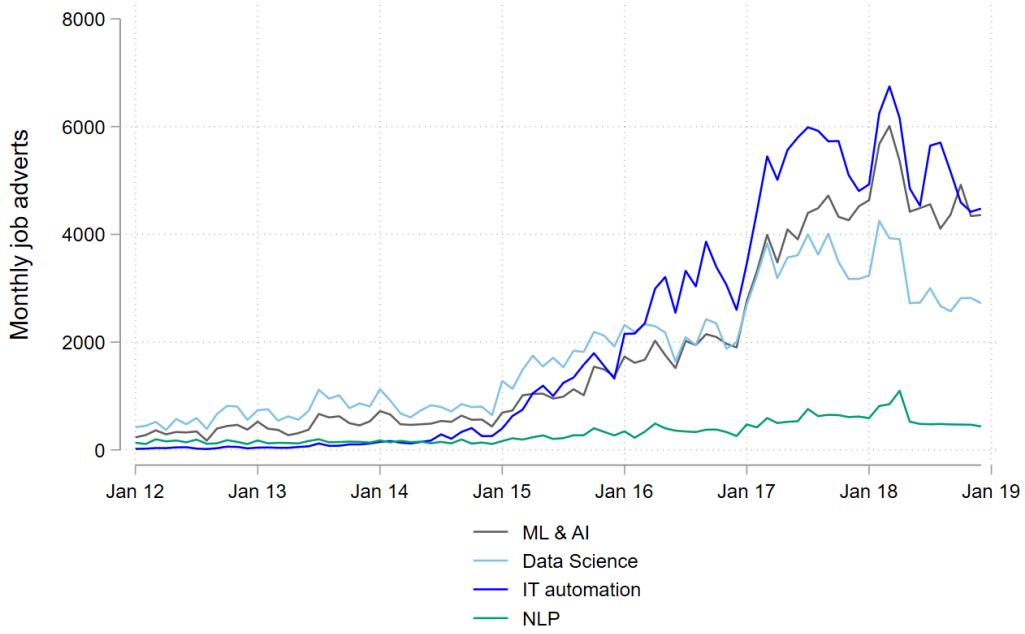
As has been noted by Seamans and Raj (2018), there is currently very little available data on AI adoption, particularly at the firm level. In the absence of data on the adoption of specific technologies, a growing body of work has started using technology-related human capital to proxy for technology adoption. For example, Rock (2019) and Benzell et al. (2019) use LinkedIn profiles to construct firm-level measures of engineering and IT talent and Harrigan et al. (2020) use the firm level employment share of 'technology workers' using French matched worker-firm data as a measure of technology adoption.

Human capital is one of the key inputs for deploying an ML system. It is well recognised that one of the primary obstacles to widespread adoption of AI is the available labour supply, with top-tier scientists earning extremely high salaries and being bought out of academic positions. It would be expected that firms wishing to implement an ML driven automation project would, at least to some degree, need to hire individuals with ML skills or experience. Alternative options would be to rely on external consultants, contract out the process to a third party software provider or retrain existing staff to develop ML skills. There are a number of reasons to believe that the dominant channel is external hiring. Work by Chui and Malhotra (2019) surveying around 2000 companies found that the primary method for sourcing AI talent and capabilities was to hire externally and that the majority of companies built their AI capabilities in house, as opposed to buying or licensing capabilities from large technology companies. Additionally, even if firms were to subcontract AI development, it would be expected that they would still require at least some related human capital in-house to oversee and manage the process.

I therefore focus on the demand for ML skills as a proxy for ML deployment. I construct measures of ML skills demand as follows. My first step is to make use of BGT's skill clusters. The job adverts collected by BGT include around 12,500 individual skills listed by employers in the UK. BGT develops skill clusters by looking at the pairwise frequency that two skills appear in the same advert and



(a) Most demanded ML & AI skills 2012-17



(b) Growth in demand over time

Figure 3.3: ML-related skills

Notes: Figure a) displays the total number of UK job adverts over the combined period of 2012-2017 including the specified skills. Figure b) displays the total number of job adverts per month including any skill in the specified skill cluster. *Source:* BGT dataset.

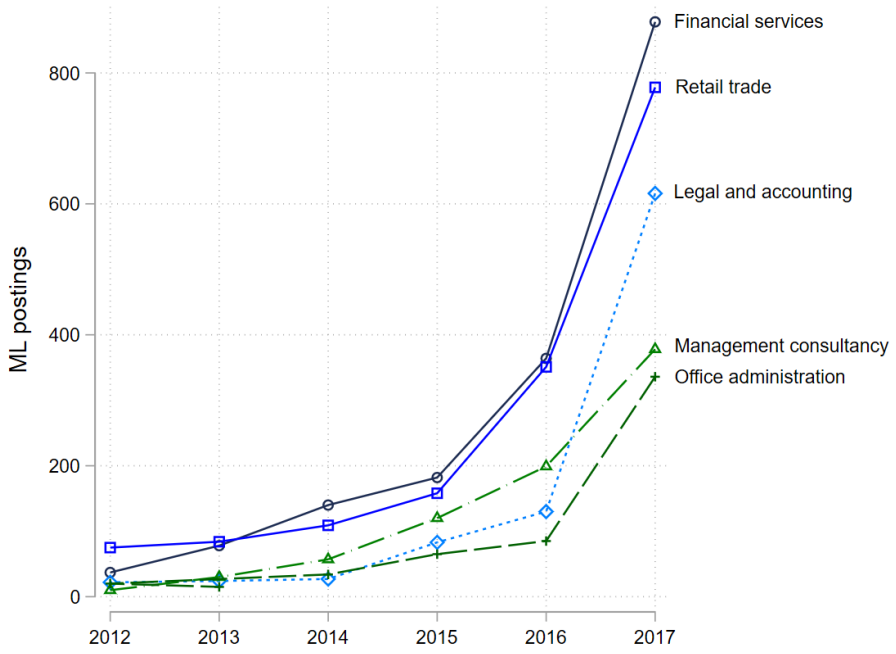
training a model to learn the extent to which skills occur in the same context. This results in 584 skill clusters.

I develop a narrow ML measure using skills from the clusters 'Machine Learning' and 'Artificial Intelligence'. I then develop a broader measure by additionally using the skills from the clusters 'NLP', 'Data Science' and 'IT automation'. I classify each post as ML-related if it includes any skill in any of these clusters. Figure 3.3(a) shows the most mentioned skills within the narrow measure clusters. The most commonly demanded skills in this group are 'Machine Learning' and 'Artificial Intelligence' themselves, followed by 'Computer Vision', 'Deep Learning' and 'Neural Networks'. Figure 3.3(b) displays the growth in job adverts including the skills within each of these five clusters used. The growth in demand for ML & AI skills and IT automation are closely linked and both grew substantially between mid 2014 and 2018. Demand for data science skills also grew substantially from mid 2014 but then slowed a bit more in mid 2017. Demand for NLP skills has been much more muted over this period.

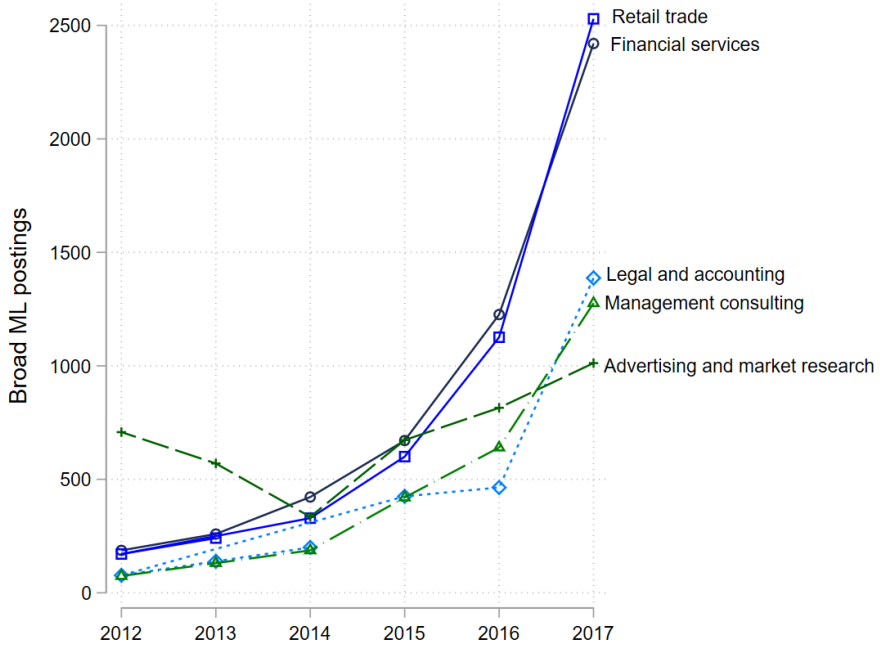
I finally take the count of all job adverts classified as ML-related within an establishment, industry or local labour market and scale it by the total number of job adverts in that establishment, industry or local labour market in the same year to develop the ML share measures. Given the focus of this paper is on ML deployment, not ML research or production, I exclude the tech sector (SIC07 codes 62 & 63) and higher education (SIC07 code 854).

Figure 3.4(a) displays the top 5 SIC 2 industries with the highest number of job adverts demanding narrowly defined ML skills over this time period. The industry with the greatest demand is financial services, followed by retail trade, legal and accounting, management consulting and office administration and business support. The industries with the next highest demand, not appearing in this graph are advertising and market research, publishing activities, other professional, scientific and technical activities, insurance, reinsurance and pension funding and scientific research and development.

Figure 3.4(b) uses the broadly defined measure of ML skills instead. The top industries are similar, except that advertising and market research is now the industry with the fifth highest demand and retail trade is now the industry with the



(a) Narrowly defined ML



(b) Broadly defined ML

Figure 3.4: Top 5 industries by demand for ML skills

Notes: These figures display the number of job adverts including ML & AI skills by SIC 2 industry in each year for the UK's top 5 industries in terms of their ML hiring. Panel (a) defines ML using the narrow measure. Panel (b) defines ML using the broad measure. *Source:* BGT data

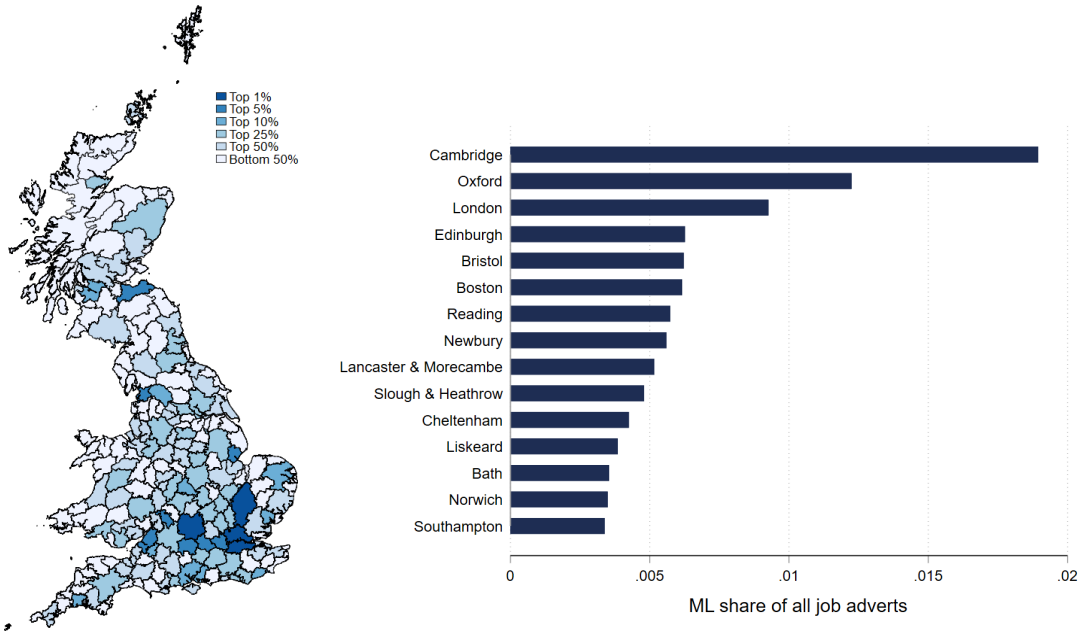


Figure 3.5: Geography of ML hiring

Notes: The map to the left displays the UK's TTWAs with the highest ML hiring intensity in 2017, using the narrowly defined measure. Darker blue colours represent greater ML hiring intensity. The bar chart to the right displays the ML hiring intensity for the 15 TTWAs with the highest intensity in 2017. *Source:* BGT data.

highest demand, followed closely by financial services. These findings are generally consistent with the GPT characteristic of AI being widely applicable across industries. Figure 3.5 displays the TTWAs with the greatest ML hiring intensity in 2017, measured as narrowly defined ML job adverts as a share of total job adverts. Cambridge, Oxford, London, Edinburgh and Bristol all have the greatest demand for ML skills.

3.5 Empirical strategy

The aim of this paper is to explore how the deployment of ML in the UK affects labour demand or services offshoring to LMI countries. The key identification challenge is that ML deployment may be driven by firm, local labour market or industry specific shocks that also affect labour demand or service sector offshoring. There is also the concern that a rise in service sector offshoring could raise productivity, resulting in greater innovation or ML deployment at the firm's headquarters. There is some evidence that this could be a possibility. Bernard et al. (2019)

document that firms tend to reorganise production in their headquarters after offshoring, with a greater focus on R&D and innovation activities.

I therefore aim to isolate improvements in ML capabilities that are exogenous to the firms adopting them. I do so by instrumenting for ML demand using exposure to advances in task-specific AI capabilities. The key idea is that advances in ML capabilities have mainly been driven by the supply side of AI research at universities or large technology firms and have been caused by factors such as expanded access to data from the use of smartphones, improved computing power or academic breakthroughs in specific fields like the use of neural networks. These technological advances were largely exogenous to firms in other sectors, such as finance, legal or retail, but made it suddenly become commercially viable to deploy ML algorithms to conduct some specific tasks. To measure the raw technology supply side of ML advances and how applicable these are to different occupations, I use the scores constructed by Felten et al. (2019), which make use of the EFF's AI Progress Metrics, mapped to the O*NET database developed by the US Department of Labor. I discuss these in detail below.

3.5.1 Measuring progress in AI capabilities

The performance of ML algorithms has improved substantially over the past decade, but they are still only applicable to a relatively narrow subset of tasks and processes. Felten et al. (2019) map specific AI capabilities, as measured and monitored by the EFF AI Progress Metrics, to the tasks most related to those capabilities and the occupations regularly performing those tasks. More details about their methodology can be found in their paper. I provide a brief description below.

The EFF AI Progress Measurement project is an open source initiative that aims to track progress on task-specific AI performance metrics across a variety of separate AI categories, such as image recognition, speech recognition or translation. For each of the categories, the EFF monitors progress in the field drawing on data from a variety of sources, including blog posts and websites focused on subfields of ML, academic literature, and review articles. The EFF aims to create the first integrated database that aggregates performance metrics across a variety of AI categories in one single place.

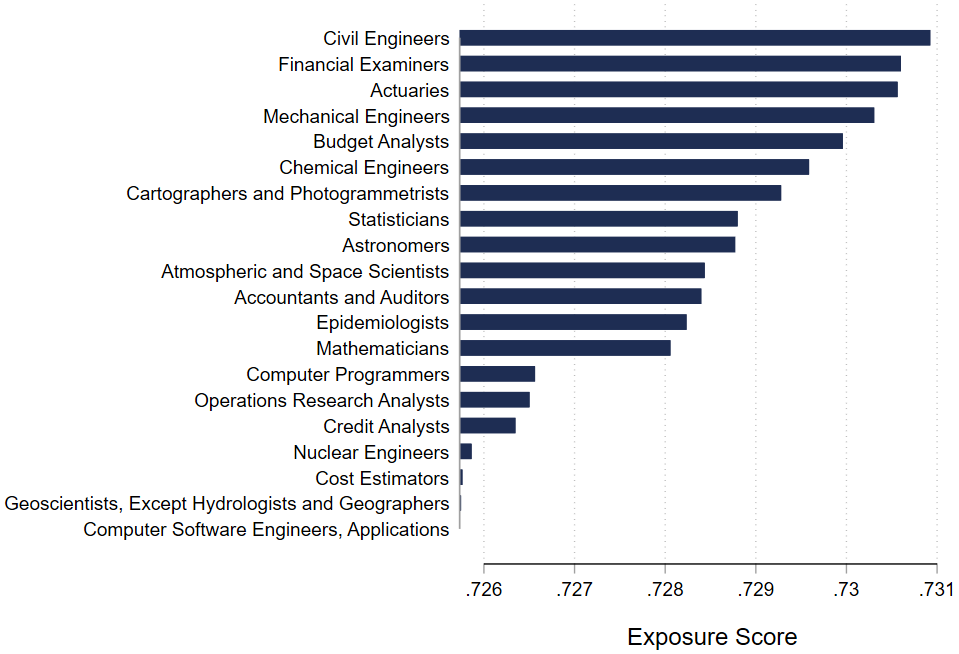


Figure 3.6: Occupations with highest exposure scores

Notes: This figure displays the 20 UK SOC 2010 codes with the highest AI exposure scores using Felten et al. (2019)'s measures, mapped to UK SOC codes.

The EFF database can be used to track the rate of change across 9 of the separate categories of metrics the EFF covers: abstract strategy games, real-time video games, image recognition, visual question answering, image generation, reading comprehension, language modeling, translation, and speech recognition. Felten et al. (2019) have integrated and aggregated all of the different metrics tracked within each category. They then map the EFF AI categories to the list of 52 abilities in the O*NET database using survey responses from a large sample of Amazon mTurk respondents. O*NET is a database that provides definitions for US occupation codes. For each of the around 1,000 occupations listed, O*NET provides information regarding personal requirements, characteristics, experience requirements and job requirements. O*NET maintains a list of 52 'abilities' and notes how important and prevalent each ability is in the relevant occupation. Examples of abilities are 'deductive reasoning' or 'speech recognition'. After mapping the EFF metrics to the O*NET abilities, Felten et al. (2019) then create aggregate AI exposure scores for each occupation using the extent to which each ability is required

for an occupation.

I use a crosswalk between O*NET occupations and UK SOC codes to obtain AI exposure measures for each UK occupation. Figure 3.6 displays the 20 occupations with the highest exposure scores. The most exposed are Civil Engineers, Financial Examiners, Actuaries, Mechanical Engineers and Budget Analysts. I then weight these measures with the establishment, local labour market or industry level BGT baseline occupation posting shares, as described below.

3.5.2 Empirical specification

The structural equation aims to investigate how long differenced changes in ML demand, measured as the ratio of ML job adverts to total job adverts, ML share$_{it}$, for unit i in year t, affect changes in the outcome variables X$_{it}$. I instrument for the change in actual ML skills demand, using the 'exposure' to advances in AI using Felten et al. (2019)'s metrics. The second and first stage, respectively, are as follows:

$$\Delta \text{X}_{it} = \Delta \text{ML share}_{it} + \text{controls}_{i,base} + \gamma_r + \gamma_j + \epsilon_{it} \tag{3.1}$$

$$\Delta \text{ML share}_{it} = \text{ML exposure}_{i,base} + \text{controls}_{i,base} + \gamma_r + \gamma_j + \text{u}_{it} \tag{3.2}$$

where controls$_{i,base}$ are controls variables in the baseline period, γ_r are region fixed effects, γ_j are industry fixed effects and ML exposure$_{i,base}$ is the firm's machine learning exposure score, measured as the baseline postings-weighted average occupational exposure score:

$$\text{ML exposure}_{i,base} = \sum_{o=1}^{O} \text{postings share}_{io}^{base} \times \text{ML exposure}_{o} \tag{3.3}$$

where postings share$_{io}^{base}$ is the postings share of occupation o in unit i in the base period as a proxy for the employment share for that occupation, and ML exposure$_o$ the exposure score for occupation o.

3.5.3 Firm level analysis for UK labour demand

For the analysis relating to the impact of ML on labour demand in the UK, I conduct the analysis at the firm, or more specifically establishment, level, with the

construction of establishment-level variables as described in Section 3.3. For this part of the analysis I take long differences between 2012 & 2013 and 2016 & 2017 to account for the fact that firms may only post adverts and recruit for a representative sample of occupations relatively infrequently. I keep only firms with at least 10 job postings for both the periods 2012 & 2013 and 2016 & 2017, which leaves a resultant sample size of 15,260 UK establishments. As a robustness check I also present the results keeping only those with at least 20 job postings.

The identifying assumption is that the baseline postings shares only affect the growth in labour demand through their impact on ML demand. Ideally I would take the shares from a year prior to any advances in AI having occurred to ensure firms or individuals hadn't already adjusted to perceived future changes in labour demand. For the analysis on offshoring it would also be preferable to use shares prior to any offshoring having occurred. Unfortunately, however, BGT only started collecting data from 2012 onwards for the UK and so 2012 is the earliest possible year I can use in this analysis. The graphs above suggested that AI demand only started increasing from 2014 and so 2012 should be prior to much ML deployment having occurred. However, there is still some possibility than firms had already started deploying ML, or offshoring in 2012. To address this concern, I include a number of baseline controls and fixed effects.

I include TTWA and SIC Division fixed effects in order to absorb any common characteristics of industries or local labour markets that could have affected both these baseline shares and subsequent growth in labour demand.[5] I also include controls for the baseline (2012 & 2013) share of postings for Administrative and Secretarial Occupations, the share for Customer Service occupations and the establishment size decile, measured in terms of baseline postings. There is a concern that in 2012 occupations had already been offshored that subsequently became feasible to automate using ML and so I control for the postings share of occupations that are typically offshored, such as Administrative and Secretarial Occupations

[5]SIC Divisions are the highest level of aggregation within the UK SIC classification. Examples are 'Wholesale and retail trade' or 'Financial and insurance activities'. I include fixed effects at this aggregation due to the limitations of the SIC classifications in the BGT data. A high proportion of postings are not classified by industry, but the classification coverage is higher and more accurate for more aggregated SIC groups.

and Customer Service occupations, in order to minimise the potential for this to affect results.

The outcome variables considered reflect total labour demand: the total number of job adverts posted, the number of non-ML job adverts, non-ML job adverts broken down by skill type and non-ML adverts posted for the most and least ML-exposed occupations, measured as those in the top and bottom 30% of the distribution in terms of exposure scores.

3.5.4 Local labour market and industry level analysis for services offshoring

To evaluate the impact of ML demand on services offshoring, it is not possible to conduct the analysis at the firm level because the ITIS services offshoring data is anonymised, so I conduct the analysis primarily at the TTWA level, with industry level analysis also for comparison. For the TTWA and industry level analysis, I take the long differences between 2012 and 2017. For the TTWA analysis I include UK NUTS1 region controls, which cover the 11 UK regions (e.g. South East, South West) and baseline controls for the 2012 share of postings for Administrative and Secretatial Occupations, the share for Customer Service occupations and additionally the public sector share of employment taken from the BRES data, which could be a strong determinant of offshoring growth of a region. For the industry level analysis I also include these same baseline controls. The key outcome variable is the total value of services offshoring, measured as imports in the categories listed in Section 3, to both LMI and HI countries. I also explore the impact on specific countries and service types.

3.6 Results

3.6.1 Firm level results for UK labour demand

I first display the first stage results for how ML exposure, constructed using the Felten et al. (2019) measures, affects growth in ML skills demand, measured in both the narrowly defined and broadly defined ways discussed in Section 4. Panel (a) of Table 3.3 displays the results for regressions of the change in the ML share of postings between 2012&13 and 2016&17 on ML exposure in 2012&13, with and

TABLE 3.3. FIRM LEVEL RESULTS: IMPACT OF ML EXPOSURE ON GROWTH IN ML SKILLS DEMAND

	Narrowly defined ML				Broadly defined ML			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Dep variable: Δ ML share of postings							
Exposure	0.056***	0.060***	0.053***	0.044***	0.18***	0.20***	0.17***	0.14***
	(0.0069)	(0.0080)	(0.0076)	(0.0084)	(0.015)	(0.016)	(0.015)	(0.016)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
R-Squared	0.0060	0.0099	0.015	0.018	0.016	0.023	0.035	0.041
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y
	(b) Dep variable: Growth in ML postings (Δ IHS(x))							
Exposure	15.8***	13.7***	11.7***	9.09***	7.38***	7.04***	6.11***	5.04***
	(2.22)	(1.88)	(1.66)	(1.88)	(0.43)	(0.41)	(0.41)	(0.41)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
R-Squared	0.0082	0.034	0.038	0.043	0.041	0.080	0.094	0.11
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y

Notes: This table displays the coefficients from the regressions of the establishment level change in two measures of ML hiring between 2012&2013 and 2016&2017 on ML exposure in 2012&2013. Panel (a) has the dependent variable of the change in ML share of postings. Panel (b) has the dependent variable of the change in the IHS transformation of the total number of ML job adverts, which is approximately equivalent to the % growth rate in ML job postings. Baseline controls are the 2012&2013 share of job adverts in Administrative and Secretarial Occupations, the share in Customer Service occupations and the establishment size decile, measured using total job postings in 2012&13. Standard errors are clustered at the parent company level. * p<0.1 ** p<0.05 ***p<0.01

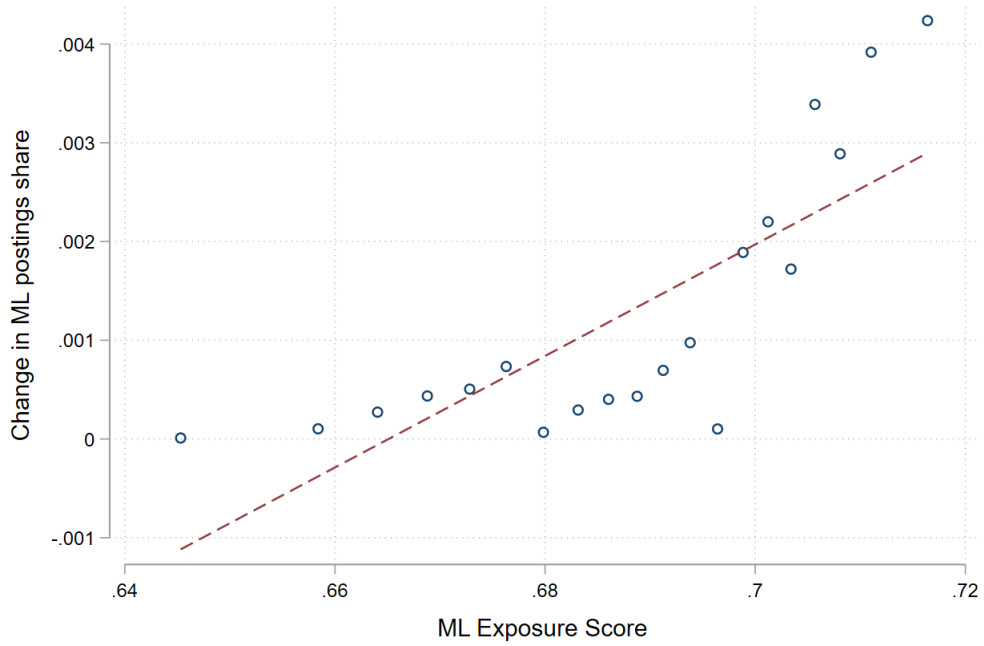


Figure 3.7: Correlation between establishment baseline ML exposure and change in demand for ML skills

Notes: This figure displays a scatter plot between the establishment baseline ML exposure scores in 2012&2013 and the change in the ML share of postings between 2012 &13 and 2016 & 17.

without controls. The results demonstrate a robust positive relationship between ML exposure and the change in demand for ML skills. Panel (b) displays the results for regressions of the change in the Inverse-Hyperbolic Sine (IHS) transformation of the total number of ML postings, between 2012&13 and 2016&17 on ML exposure in 2012&13. I take the IHS transformation because there are many establishments with zero ML postings and so taking logs would lead to the loss of a substantial proportion of the sample. This results in a metric which is an approximation of the % growth in ML postings. With this outcome variable I also find a robust positive relationship. These results generally suggest that the baseline exposure to advances in ML is a strong predictor for growth in the demand for ML skills, as measured through job adverts, suggesting this would make a strong IV.

Figure 3.7 also graphs the relationship between establishment level baseline exposure scores and the change in the ML share of postings, displaying the positive, although slightly non-linear, relationship. Table 3.4 shows the second stage of these IV regressions with the outcome variables of the change in the log of non-ML

TABLE 3.4. FIRM LEVEL IV RESULTS: IMPACT OF ML DEMAND ON NON-ML & TOTAL POSTINGS

	Narrowly defined ML				Broadly defined ML			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Dep variable: Growth in non-ML postings (Δ log postings)							
ΔML share	-0.70	55.6***	54.7***	122.0***	-0.22	16.7***	16.8***	37.3***
	(13.8)	(16.1)	(16.7)	(34.1)	(4.21)	(4.51)	(4.82)	(8.50)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
First stage F stat	67.7	54.7	47.2	24.5	158.6	150.1	117.8	66.0
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y
	(b) Dep variable: Growth in total postings (Δ log postings)							
ΔML share	2.83	59.2***	58.2***	125.6***	0.87	17.7***	17.9***	38.4***
	(13.7)	(16.3)	(16.8)	(34.3)	(4.21)	(4.50)	(4.82)	(8.50)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
First stage F stat	67.7	54.7	47.2	24.5	158.6	150.1	117.8	66.0
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y

Notes: This table displays the coefficients from the second stage regressions of the establishment level change in non-ML postings or total postings on the change in the ML share of postings between 2012&2013 and 2016&2017. The instrument used is the ML exposure in 2012/13. Baseline controls are the 2012&2013 share of job adverts in Administrative and Secretarial Occupations, the share in Customer Service occupations and the establishment size decile, measured using total job postings in 2012&13. Standard errors are clustered at the parent company level. * p<0.1 ** p<0.05 ***p<0.01

postings and the change in the log of total postings. For columns 1-4, non-ML postings are non-narrowly defined ML postings and for columns 5-8 non-ML postings are non-broadly defined ML postings. The first stage F stats are generally high in all specifications, providing further evidence that the exposure scores are a strong predictor of growth in ML postings.

These regressions suggest that the growth in demand for ML skills hasn't led to a decrease in labour demand or demand for non-ML positions and, on the contrary, there is evidence of a positive impact on both growth in non-ML postings and total postings. The significance and magnitudes of the coefficients increase substantially only after the baseline controls are included. When added in iteratively, it is the control for the firm size decile that has the biggest impact on the

TABLE 3.5. FIRM LEVEL IV RESULTS: IMPACT OF ML DEMAND ON NON-ML POSTINGS BY SKILL TYPE

	Narrowly defined ML				Broadly defined ML			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Dep variable: Growth in non-ML higher-skilled postings (Δ IHS postings)							
ΔML share	-54.7***	9.13	-2.85	11.8	-17.7***	1.80	-1.78	2.70
	(18.4)	(17.1)	(17.5)	(29.8)	(5.42)	(5.11)	(5.39)	(9.09)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
First stage F stat	67.7	54.7	47.2	24.5	158.6	150.1	117.8	66.0
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y
	(b) Dep variable: Growth in non-ML lower-skilled postings (Δ IHS postings)							
ΔML share	75.2***	125.2***	133.9***	279.5***	11.0*	25.2***	28.8***	72.5***
	(20.3)	(25.6)	(28.6)	(65.4)	(5.73)	(6.21)	(7.22)	(13.7)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
First stage F stat	67.7	54.7	47.2	24.5	158.6	150.1	117.8	66.0
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y

Notes: This table displays the coefficients from the second stage regressions of the establishment level change in non-ML higher-skilled postings or non-ML lower-skilled postings on the change in the ML share of postings between 2012&2013 and 2016&2017. The instrument used is the ML exposure in 2012/13. Baseline controls are the 2012&2013 share of job adverts in Administrative and Secretarial Occupations, the share in Customer Service occupations and the establishment size decile, measured using total job postings in 2012&13. Standard errors are clustered at the parent company level. * p<0.1 ** p<0.05 ***p<0.01

coefficient. This suggests that, for firms of a similar size, an increase in ML demand increases demand for other occupations, but across firms of different sizes this relationship doesn't hold. This could imply that the relationship varies across the firm size distribution. Controlling for the firm's TTWA has little additional effect on the coefficients, but controlling for the SIC industry division leads the coefficients to approximately double, suggesting that for firms of a similar size within the same industry, growth in ML demand has an even greater effect on the growth in postings for non-ML vacancies.

In terms of magnitudes, looking at the narrow definition and the specification in column 3, a one standard deviation increase in the growth in the ML share of postings (0.01) leads to around a 55% increase in non-ML postings and a 58% increase in total postings over this five year period. For the broad definition and specification in column 7, a one standard deviation increase in the growth in the

ML share of postings (0.03) leads to around a 50% increase in non-ML postings and a 54% increase in total postings.

In Table 3.5 I next look at this impact on non-ML job adverts by skill type. Panel (a) studies the impact on higher skilled, non-ML postings, with this skill classification defined in the Data section. Without any controls added, the increase in the ML share of postings seems to have a negative effect on the growth in higher skilled postings. However, once the baseline controls are added these coefficients lose significance and for the further three specifications there appears to be no statistically significant impact on higher skilled postings. Panel (b) studies the impact on lower-skilled non-ML postings, demonstrating that the growth in non-ML postings appears to be driven by increased demand for lower-skilled roles. This is particularly the case for the narrowly defined ML measure, which has strongly significant coefficients in all specifications. These results could suggest that the increase in demand for ML skills leads to an increase in demand for complementary occupations in the UK and these complementary occupations tend to be lower skilled, such as data taggers. Alternatively, these results could reflect scale effects.

I further explore whether this impact on non-ML occupations is influenced by whether or not occupations are automatable by ML. Table 3.6 displays the second stage results for the impact of growth in ML skills demand on the growth in non-ML job postings for the most and least ML-exposed occupations, measured using the Felten et al. (2019) measures, taking the 30% cutoff for both. Panel (a) looks at the impact of ML demand on the posting of adverts for the top 30% most exposed occupations, excluding ML postings. The results are negative and very statistically significant in all specifications, suggesting that an increase in ML skills demand reduces labour demand for occupations that are substitutable with ML. Taking specification (3), for example, a one standard deviation increase in the establishment's growth in the ML share of postings (0.01) leads to an approximately 117% decline in postings for the most exposed occupations.

Panel (b) looks at the impact on the least exposed non-ML postings. The results are positive and significant in most of the specifications, but less so for the broad ML measure. These results suggest that an increase in ML skills demand seems to increase labour demand for non-ML exposed occupations, but some of these occu-

TABLE 3.6. FIRM LEVEL IV RESULTS: IMPACT OF ML DEMAND ON MOST & LEAST EXPOSED OCCUPATIONS

	Narrowly defined ML				Broadly defined ML			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Dep variable: Growth in most exposed non-ML postings (Δ IHS(x))							
ΔML share	-121.9***	-81.9***	-116.6***	-180.3***	-38.3***	-25.5***	-36.8***	-56.0***
	(21.7)	(19.7)	(24.6)	(47.3)	(5.64)	(5.24)	(6.30)	(11.5)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
First stage F stat	67.7	54.7	47.2	24.5	158.6	150.1	117.8	66.0
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y
	(b) Dep variable: Growth in least exposed non-ML postings (Δ IHS(x))							
ΔML share	39.7*	99.2***	135.5***	357.1***	-13.5**	6.00	17.6**	85.0***
	(20.3)	(24.7)	(31.3)	(84.0)	(5.78)	(5.85)	(7.25)	(16.8)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
First stage F stat	67.7	54.7	47.2	24.5	158.6	150.1	117.8	66.0
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y

Notes: This table displays the coefficients from the second stage regressions of the establishment level change in posting of non-ML job adverts in the 30% most and least ML-exposed occupations on the change in the measures of ML demand between 2012&2013 and 2016&2017. The instrument used is the ML exposure in 2012/13. Baseline controls are the 2012&2013 share of job adverts in Administrative and Secretarial Occupations, the share in Customer Service occupations and the establishment size decile, measured using total job postings in 2012&13. Standard errors are clustered at the parent company level. * p<0.1 ** p<0.05 ***p<0.01

TABLE 3.7. TTWA IV RESULTS: IMPACT OF ML DEMAND ON SERVICES OFFSHORING

	Narrowly defined ML			Broadly defined ML		
	(1)	(2)	(3)	(4)	(5)	(6)
	(a) Dep variable: Δ Log offshoring to LMI countries					
Δ ML share	850.6***	925.2***	1026.4***	256.5***	262.3***	273.9***
	(251.7)	(285.8)	(316.4)	(68.5)	(72.8)	(73.7)
First stage coefficient	0.093***	0.090***	0.091***	0.31***	0.32***	0.34***
	(0.020)	(0.021)	(0.022)	(0.052)	(0.058)	(0.060)
First stage F stat	22.9	18.9	18.0	35.2	30.2	32.8
Observations	213	213	213	213	213	213
NUTS1 FE	N	Y	Y	N	Y	Y
Baseline controls	N	N	Y	N	N	Y
	(b) Dep variable:Δ Log offshoring to HI countries					
Δ ML share	193.0	256.5*	206.6	58.3	72.6*	55.1
	(123.7)	(137.3)	(145.3)	(36.8)	(37.1)	(37.5)
First stage coefficient	0.094***	0.090***	0.092***	0.31***	0.32***	0.35***
	(0.020)	(0.021)	(0.022)	(0.053)	(0.058)	(0.060)
First stage F stat	22.7	18.8	18.1	34.7	30.2	32.9
Observations	212	212	212	212	212	212
NUTS1 FE	N	Y	Y	N	Y	Y
Baseline controls	N	N	Y	N	N	Y

Notes: This table displays the coefficients from the 2SLS regressions of the TTWA level change in the log of offshoring between 2012 and 2017 on the change in the ML share of job adverts between 2012 and 2017. The instrument used is the TTWA ML exposure in 2012&2013. Panel (a) has the dependent variable of the change in the log of total offshoring to LMI countries. Panel (b) has the dependent variable of the change in the log of total offshoring to HI countries. Baseline controls are the 2012 share of job adverts in Administrative and Secretarial Occupations, the share in Customer Service occupations and the public sector share of employment in 2012. Robust standard errors in parentheses. * p<0.1 ** p<0.05 ***p<0.01

pations are perhaps related to data science or IT automation and hence the impact on least exposed non-broad ML postings is lower. In terms of magnitudes, for specification (3), a one standard deviation increase in the establishment's growth in the ML share of postings (0.01) leads to an approximately 136% increase in postings for the least exposed occupations, quite a bit higher than the negative impact for exposed occupations. These results generally suggest that there does appear to be some degree of substitution between ML-skills and ML exposed occupations, although ML deployment also increases demand for non-exposed occupations.

3.6.2 Local labour market results for services offshoring

Table 3.7 displays the TTWA level results for services offshoring. At the TTWA level, the first stage between TTWA level ML exposure in 2012 and the change in the ML share of postings between 2012/13 and 2017 remains strong, particularly for the measures of broadly defined ML hiring. Panel (a) looks at the impact of local labour market growth in demand for ML skills on the growth in offshoring to LMI countries between 2012 and 2017, with and without controls. The coefficients are positive and strongly significant, demonstrating that growth in ML demand has a positive effect on growth in services offshoring to LMI countries. In terms of magnitude, taking column (3) as the baseline, a one standard deviation increase in the TTWA level change in the ML share of postings (0.001) leads to an approximately 103% increase, or doubling, of the value of services offshoring to LMI countries between 2012 and 2017.

Panel (b) looks at the impact of local labour market growth in demand for ML skills on the growth in offshoring to HI countries between 2012 and 2017, with and without controls. The coefficients only have limited significance in two specifications, suggesting that there is insignificant evidence of ML demand having an impact on offshoring to HI countries.

Table 3.8 displays the results for the breakdown of impact by country for the UK's 4 largest services import partners in the LMI category and aggregated service types, with the details on these groups provided in Appendix 3.A. Panel (a) shows that there is a particularly strong impact of growth in ML demand on offshoring to India, with some evidence of an increase to South Africa. Brazil and China, on the other hand, have not been affected. This holds when looking at either narrowly defined or broadly defined ML. In terms of service types, panel (b) shows that the increase in offshoring to LMI countries appears to have been concentrated in the categories of Business Services and ICT-related services (telecommunications services, computer services and information services), while Technical Services and Services between related parties n.e.i. have not been affected.

TABLE 3.8. TTWA IV RESULTS: IMPACT OF ML DEMAND ON OFFSHORING BY LMI COUNTRY & SERVICE TYPE

	Narrowly defined ML				Broadly defined ML			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Dep variable: Δ Log service offshoring from partner country							
	India	China	Brazil	SA	India	China	Brazil	SA
ΔML share	851.3***	-108.0	546.5	682.2**	228.8***	-30.4	157.7	190.8**
	(247.4)	(284.5)	(380.2)	(295.0)	(58.0)	(79.7)	(108.6)	(80.2)
First stage F stat	19.2	21.9	25.1	22.9	33.7	32.0	31.7	33.7
Observations	207	188	174	184	207	188	174	184
NUTS FE	Y	Y	Y	Y	Y	Y	Y	Y
Baseline controls	Y	Y	Y	Y	Y	Y	Y	Y
	(b) Dep variable: Δ Log service offshoring by service type							
	Business	ICT	Technical	Other Related	Business	ICT	Technical	Other Related
ΔML share	1225.3***	151.0***	71.5	115.5	207.9**	326.9***	20.1	30.6
	(353.5)	(365.3)	(368.2)	(308.7)	(88.3)	(84.1)	(103.5)	(81.3)
First stage F stat	18.7	18.0	25.8	19.8	34.4	32.8	35.3	33.9
Observations	209	213	148	161	209	213	148	161
NUTS FE	Y	Y	Y	Y	Y	Y	Y	Y
Baseline controls	Y	Y	Y	Y	Y	Y	Y	Y

Notes: This table displays the coefficients from the 2SLS regressions of the TTWA level change in the log of offshoring to LMI countries between 2012 and 2017 on the change in the ML share of job adverts between 2012 and 2017. The instrument used is the TTWA ML exposure in 2012&2013. Panel (a) has the dependent variable of the change in the log of total offshoring to LMI countries by country, for the four largest LMI import partners. Panel (b) has the dependent variable of the change in the log of total offshoring to LMI countries by service type. Baseline controls are the 2012 share of job adverts in Administrative and Secretarial Occupations, the share in Customer Service occupations and the public sector share of employment in 2012. Robust standard errors in parentheses. * p<0.1 ** p<0.05 ***p<0.01

TABLE 3.9. INDUSTRY IV RESULTS: IMPACT OF ML DEMAND ON OFFSHORING

	Narrowly defined ML		Broadly defined ML	
	(1)	(2)	(3)	(4)
	(a) Dep variable: Δ IHS offshoring to LMI countries			
Δ ML share	124.3***	137.5***	53.6***	56.2***
	(66.1)	(65.2)	(25.8)	(24.15)
First stage coefficient	0.083***	0.084***	0.19***	0.19*** *
	(0.021)	(0.021)	(0.044)	(0.045)
First stage F stat	15.9	20.8	18.6	22.5
Observations	246	246	246	246
Baseline controls	N	Y	N	Y
	(b) Dep variable:Δ IHS offshoring to HI countries			
Δ ML share	198.1	199.4*	85.5**	81.5**
	(99.6)	(91.4)	(40.0)	(35.1)
First stage coefficient	0.083***	0.084***	0.19***	0.19***
	(0.021)	(0.021)	(0.044)	(0.045)
First stage F stat	15.9	20.8	18.6	22.5
Observations	246	246	246	246
Baseline controls	N	Y	N	Y

Notes: This table displays the coefficients from the 2SLS regressions of the SIC 3 digit industry level change in the IHS of offshoring between 2012 and 2017 on the change in the ML share of job adverts between 2012 and 2017. The instrument used is the SIC 3 digit industry level ML exposure in 2012&2013. Panel (a) has the dependent variable of the change in the IHS of total offshoring to LMI countries. Panel (b) has the dependent variable of the change in the IHS of total offshoring to HI countries. Baseline controls are the 2012 share of job adverts in Administrative and Secretarial Occupations, the share in Customer Service occupations and the public sector share of employment in 2012. Standard errors are clustered by SIC 2 digit industry. * p<0.1 ** p<0.05 ***p<0.01

TABLE 3.10. INDUSTRY IV RESULTS: IMPACT OF ML DEMAND ON OFFSHORING BY COUNTRY & SERVICE TYPE

	Narrowly defined ML				Broadly defined ML			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Dep variable: Δ IHS service offshoring from partner country							
	India	China	Brazil	SA	India	China	Brazil	SA
ΔML share	94.7***	-0.44	66.1*	52.1*	38.7**	-0.18	27.0*	21.3*
	(42.7)	(23.7)	(40.1)	(29.9)	(16.6)	(9.70)	(16.1)	(11.5)
First stage F stat	20.8	20.8	20.8	20.8	22.5	22.5	22.5	22.5
Observations	246	246	246	246	246	246	246	246
Baseline controls	Y	Y	Y	Y	Y	Y	Y	Y
	(b) Dep variable: Δ IHS service offshoring by service type							
	Business	ICT	Technical	Other Related	Business	ICT	Technical	Other Related
ΔML share	146.3***	-4.52***	36.7	-43.6	59.8***	-1.85	15.0	-17.8
	(56.2)	(31.6)	(43.3)	(44.8)	(21.3)	(13.0)	(17.1)	(18.6)
First stage F stat	20.8	20.8	20.8	20.8	22.5	22.5	22.5	22.5
Observations	246	246	246	246	246	246	246	246
Baseline controls	Y	Y	Y	Y	Y	Y	Y	Y

Notes: This table displays the coefficients from the 2SLS regressions of the SIC 3 digit industry level change in the IHS of offshoring between 2012 and 2017 on the change in the ML share of job adverts between 2012 and 2017. The instrument used is the SIC 3 digit industry level ML exposure in 2012&2013. Panel (a) has the dependent variable of the change in the IHS of total offshoring to LMI countries by country for the top LMI import partners. Panel (b) has the dependent variable of the change in the IHS of total offshoring to LMI countries by service product. Baseline controls are the 2012 share of job adverts in Administrative and Secretarial Occupations, the share in Customer Service occupations and the public sector share of employment in 2012. Standard errors are clustered by SIC 2 digit industry. * p<0.1 ** p<0.05 ***p<0.01

3.6.3 Industry level results

Table 3.9 repeats the same analysis as for TTWAs but using SIC 3 digit industries as the unit of analysis instead. For these industry level regressions I take the IHS transformation of the value of offshoring because there is a high proportion of industries with zero offshoring. At the industry level, there is also a strong first stage and evidence of a positive impact of growth in ML demand on services offshoring to LMI countries, but there is now also a positive impact on HI countries as well. Table 3.10 displays the breakdown by import partner country and service type. I again find a strong positive impact particularly for India, with some evidence of a positive impact for South Africa as well, and now some evidence of a positive impact for Brazil. I now find primarily a strong impact for business services, with the impact for ICT services now being insignificant or negative.

3.7 Robustness

3.7.1 Firm level robustness checks

Replacing ML share with ML postings

In the second stage analysis I used the ML share of postings as the primary measure of the change in ML demand. I did so to mitigate the potential issue of some firms, industries or areas of the economy advertising a higher proportion of jobs online. This therefore assumed that within a firm, the share of jobs posted online would be roughly similar between ML and non-ML jobs. One concern could be that this is not the case and that smaller or growing firms could be more likely to hire through informal networks, but post online for only a subset of roles, such as those that are niche or technology related roles, such as for ML expertise. This could result in an overstatement of the growth in the ML share for these companies, while large firms may be more likely to post a more representative sample of roles online. If smaller firms are also conducting more expansion related hiring, then this could result in an inflated estimate for the impact of ML on labour demand.

I therefore also repeat the main establishment level results using the change in the total number of ML postings, transformed by IHS. I demonstrate in Table 3.11 that doing so does not change the conclusions about the impact of ML demand on postings for non-ML roles. In terms of magnitude, taking the narrow definition and Column 3, a one sd (3.2) increase in the change in the IHS transformation of ML postings leads to an approximately 80% increase in non-ML postings and an approximately 83% increase in total postings over this five year period, slightly higher than the magnitude of the coefficient using the ML share, suggesting that using the ML share is not leading to an overestimate of the effect, but perhaps vice versa.

Increasing minimum threshold to 20 postings

The main establishment level sample consisted of establishments with a minimum of 10 postings for both the periods of 2012&13 and 2016&17. An additional concern

TABLE 3.11. FIRM LEVEL IV RESULTS: REPLACING ΔML SHARE WITH ΔIHS(ML POSTINGS)

	Narrowly defined ML				Broadly defined ML			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Dep variable: Growth in non-ML postings (Δ log(postings))							
ΔIHS(ml postings)	-0.0025	0.24***	0.25***	0.58***	-0.0054	0.47***	0.46***	1.00***
	(0.049)	(0.069)	(0.071)	(0.17)	(0.11)	(0.12)	(0.12)	(0.20)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
First stage F stat	50.6	51.1	46.5	20.6	291.7	278.4	213.2	135.7
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y
	(b) Dep variable: Growth in total postings (Δ log(postings))							
ΔIHS(ml postings)	0.010	0.26***	0.26***	0.59***	0.022	0.50***	0.49***	1.03***
	(0.049)	(0.070)	(0.072)	(0.17)	(0.11)	(0.12)	(0.12)	(0.20)
Observations	15260	15260	15251	15251	15260	15260	15251	15251
First stage F stat	50.6	51.1	46.5	20.6	291.7	278.4	213.2	135.7
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y

Notes: This table replicates Table 3.4 using ΔIHS(ml postings) instead of ΔML share.

could be that this threshold is too low and the resultant sample includes establishments that are not posting a representative sample of their positions online or are not postings consistently with the same name. I therefore also evaluate whether using a higher minimum threshold of 20 alters the results. This reduces the sample size to 7,730 establishments. Table 3.12 displays the results for non-ML postings and total postings. When conducting the same analysis using this reduced sample of firms with more postings, the coefficients in specifications 1, 2, 3 and 5, 6, 7 are all either lower in magnitude or lower in significance, but for columns and 8, when controlling for SIC division, the coefficients are very similar. These results imply that when the effect is estimated between firms in the same industry, including smaller firms does not matter, but when estimating the effect across industries, it does matter. In general, however, the results remain qualitatively similar in that ML seems to have a weakly positive impact on labour demand, once controlling for other potentially confounding factors.

TABLE 3.12. FIRM LEVEL IV RESULTS: INCREASING THRESHOLD TO 20

	Narrowly defined ML				Broadly defined ML			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Dep variable: Growth in non-ML postings (Δ log(postings))							
ΔML share	-12.8	26.9*	35.8**	125.8***	-4.09	8.50*	11.6**	37.0***
	(12.6)	(14.3)	(14.5)	(36.5)	(4.00)	(4.50)	(4.64)	(9.16)
Observations	7752	7752	7730	7730	7752	7752	7730	7730
First stage F stat	82.6	73.6	60.8	22.4	139.1	132.5	111.2	53.6
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y
	(b) Dep variable: Growth in total postings (Δ log(postings))							
ΔML share	-9.47	30.3**	39.1***	129.4***	-3.01	9.58**	12.6***	38.1***
	(12.6)	(14.3)	(14.5)	(36.8)	(4.00)	(4.50)	(4.64)	(9.16)
Observations	7752	7752	7730	7730	7752	7752	7730	7730
First stage F stat	82.6	73.6	60.8	22.4	139.1	132.5	111.2	53.6
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y

Notes: This table replicates Table 3.4 reducing the sample to only establishments with 20 or more postings in 2012/13 and 2016/17.

Changes before the Brexit referendum

Another concern is that the Brexit vote, which occurred in June 2016, falls at the end of the period of analysis and so may have affected postings in the first or second half of 2016. I have tried to account for this by taking a long difference approach and taking averages across years. In Chapter 2 of this thesis we showed that the threat of trade barriers on professional services imports caused by the Brexit vote had a negative effect on job postings, particularly after 2018. While taking the long difference up until 2017 should therefore only have a potentially muted effect on job adverts, it remains a possibility. I therefore also repeat the establishment level results for the period of 2012-2015. Table 3.13 displays the results. Generally, for this time period the first stage is substantially weaker, which could reflect the fact that ML deployment really started to take off after 2015, so these last years are important for the analysis. The second stage results are weaker but reflect a similar pattern: there is some weak evidence of an increase in non-ML and total postings, particularly when comparing firms within the same SIC Division.

TABLE 3.13. FIRM LEVEL IV RESULTS: 2012-2015

	Narrowly defined ML				Broadly defined ML			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(a) Dep variable: Growth in non-ML postings (Δ log(postings))							
ΔML share	294.5*	189.7	258.1	623.6**	72.0**	45.5	59.5	136.8**
	(154.0)	(136.3)	(189.9)	(305.2)	(36.4)	(32.0)	(42.4)	(59.2)
Observations	10572	10572	10553	10553	10572	10572	10553	10553
First stage F stat	24.7	22.2	16.5	10.9	50.0	50.2	40.7	25.7
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y
	(b) Dep variable: Growth in total postings (Δ log(postings))							
ΔML share	295.4*	190.3	258.9	624.4**	72.2**	45.6	59.7	137.0**
	(154.0)	(136.4)	(189.9)	(305.3)	(36.4)	(32.0)	(42.4)	(59.2)
Observations	10572	10572	10553	10553	10572	10572	10553	10553
First stage F stat	24.7	22.2	16.5	10.9	50.0	50.2	40.7	25.7
Baseline controls	N	Y	Y	Y	N	Y	Y	Y
TTWA FE	N	N	Y	Y	N	N	Y	Y
SIC Division FE	N	N	N	Y	N	N	N	Y

Notes: This table replicates Table 3.4 but with long differences over the time period of 2012-2015.

Remaining threats to identification

As discussed in section 3.5, the identifying assumption for these regressions is that the 2012 & 13 exposure scores affect the long differenced outcome variables only through their impact on the growth in ML skills demand. There are a few remaining threats to identification when studying UK labour demand that have not been possible to investigate with this data. The first is that firms do not attempt to hire ML experts, but train their existing staff, outsource or contract out the development of an ML system. In this scenario, the true relationship between exposure and ML deployment would be underestimated by my specification and the labour demand effects of ML would be attributed to a low level of growth in ML. I provided anecdotal evidence in Section 3.5 that firms are more likely to hire than outsource, and the strength of the first stage relationship suggests that exposed firms are indeed posting adverts demanding ML skills, but this remains a potential issue. For the TTWA analysis I address this problem as discussed below, making use of data from the ABS, but for the establishment level this data is not available.

TABLE 3.14. TTWA LEVEL IV RESULTS: REPLACING ML SHARE WITH IHS(ML POSTINGS)

	Narrowly defined ML			Broadly defined ML		
	(1)	(2)	(3)	(4)	(5)	(6)
	(a) Dep variable: Δ Log offshoring to LMI countries					
Δ IHS(ML postings)	0.82***	0.93***	1.07***	0.91***	1.00***	1.09***
	(0.2)	(0.24)	(0.28)	(0.23)	(0.27)	(0.29)
First stage F stat	75.5	37.8	31.1	48.5	30.2	28.1
Observations	213	213	213	213	213	213
NUTS1 FE	N	Y	Y	N	Y	Y
Baseline controls	N	N	Y	N	N	Y
	(B) Dep variable:Δ Log offshoring to HI countries					
Δ IHS(ML postings)	0.19	0.26**	0.22	0.21	0.28**	0.22
	(0.12)	(0.12)	(0.14)	(0.13)	(0.14)	(0.15)
First stage F stat	74.1	37.7	31.2	46.7	29.8	28.4
Observations	212	212	212	212	212	212
NUTS1 FE	N	Y	Y	N	Y	Y
Baseline controls	N	N	Y	N	N	Y

Notes: This table replicates Tables 3.7 but replaces Δ ML share with Δ IHS(ML postings).

Another remaining possibility is that, even when pooling the data for 2012 and 2013, the baseline postings shares do not reflect a representative sample of the occupations employed in a firm. This could particularly be the case for new, small or expanding firms. If the firm has just undertaken a lot of new hiring for one expanding department or business function, it might be subsequently less likely to advertise for these roles again in the following years. Or vice versa, it may be attempting to expand one department or business function in light of new opportunities, so has a high share of certain occupations in its baseline postings sample and then high growth in those same occupations. It is unclear which direction this would affect the results, but this possibility could be particularly problematic for the regressions for the posting of most or least exposed occupations. Further work could attempt to extend this analysis using additional data on the actual baseline employment composition of establishments, as opposed to postings composition, to evaluate if this is leading to biased results.

A final potential threat could be if certain occupations have already been offshored that are highly exposed to ML. Although I control for the shares of adminis-

trative and secretarial and customer service occupations, there may be other occupations that have been offshored that are ML exposed. A firm may then have a low exposure score, but a high demand for ML to automate its business functions that have previously been offshored, again potentially leading to an underestimate of the relationship between exposure scores and ML demand. Without the firm level component of the ITIS data matched to the BGT data, it is hard to tell is this is what is indeed happening, but if it is, it would be expected to dampen the results, so should be less of a concern.

3.7.2 Local labour market -level robustness checks

Replacing ML share with ML postings

As for the firm level analysis, I also repeat the baseline results for TTWAs instead using the change in the IHS transformation of the total number of ML postings as the key measure of ML demand. For TTWAs there is also the concern that for more rural or smaller TTWAs, ML job adverts may be over represented relative to other occupations, leading to overestimates of the growth in the ML share for these firms. Table 3.14 displays the results, showing that, as with the firm level results, changing this measure does not dampen the results or qualitatively change the conclusions about offshoring.

Controlling for computer service purchases

A final key concern already mentioned is that firms do not post any job adverts including ML skills but instead hire contractors or consultants. For the TTWA analysis I attempt to mitigate this issue by adding a baseline control variable constructed from the ABS data for the 'Cost of Computer and Related Services Purchased'. I aggregate up the value of this variable for local units to TTWAs and construct a baseline control for TTWA Computer Service intensity by weighting by TTWA level employment. This allows me to compare the impact of ML demand for TTWAs already conducting a similar amount of computer service related subcontracting. Table 3.15 displays the results including this additional control. The first stage loses significance slightly with this control, but there remains a strong

TABLE 3.15. TTWA LEVEL IV RESULTS CONTROLLING FOR COMPUTER SERVICE EXPENDITURE

	Narrowly defined ML			Broadly defined ML		
	(1)	(2)	(3)	(4)	(5)	(6)
	(a) Dep variable: Δ Log offshoring to LMI countries					
Δ ML share	800.1***	896.6***	1001.1***	226.3***	255.0***	265.6***
	(264.3)	(297.8)	(328.7)	(73.2)	(77.6)	(77.3)
First stage F stat	20.8	17.0	16.2	30.3	25.9	28.5
Observations	213	213	213	213	213	213
NUTS1 FE	N	Y	Y	N	Y	Y
Baseline controls	N	N	Y	N	N	Y
	(B) Dep variable:Δ Log offshoring to HI countries					
Δ ML share	198.5	254.0*	192.8	60.2	72.1*	51.1
	(140.2)	(154.0)	(160.8)	(42.1)	(42.1)	(41.5)
First stage F stat	20.6	16.9	16.2	28.7	25.9	28.6
Observations	212	212	212	212	212	212
NUTS1 FE	N	Y	Y	N	Y	Y
Baseline controls	N	N	Y	N	N	Y

Notes: This table replicates Table 3.7, adding an additional control for the intensity of computer service purchases in 2012, measured as computer service purchases divided by employment.

impact for LMI countries, with coefficients of a similar magnitude, and an insignificant, or weakly significant, effect for HI countries.

Remaining threats to identification

For the analysis on offshoring, there are still a few remaining threats to identification. For TTWAs, one concern could be that TTWAs that are more exposed to ML are also those with a higher baseline share of highly educated workers, such as civil engineers or financial examiners, and these are also areas more likely to have a high concentration of large, productive firms that are more likely to have seen a growth in offshoring during this period. Indeed, London is the most ML exposed TTWA, followed by Edinburgh and Bristol. I try to address this concern by including NUTS1 fixed effects, which mean that the analysis compares differentially exposed TTWAs within the same region, for example, those in the South East of England and through controlling for the baseline occupational shares, that could be expected to affect offshoring growth. However, the possibility that there

are further factors that affected offshoring growth and are correlated with ML exposure, remains. Further work could attempt to evaluate this relationship using more granular data, such as at the firm level, where it would be possible to control for moe of these factors.

A final concern could be around the offshoring of ML deployment. If firms are offshoring the task of designing or implementing ML systems, then we might witness a high ML exposure score, along with low ML demand but high offshoring. Once again, this could be expected to result in an underestimate of the true relationship between exposure and ML deployment and should dampen the first stage. It should also lead to estimation of the relationship between ML demand in the UK and offshoring that is biased towards zero, again serving to only dampen the results, which is less of a concern.

3.8 Conclusions

This paper uses data on the majority of all job adverts posted online in the UK between 2012 and 2017 to develop measures for the demand for ML skills in UK firms, local labour markets and industries. I show that the posting of job adverts that demand ML skills has risen rapidly between 2012 and 2017 and there has been an increase in demand across industries. The UK local labour markets that have experienced the greatest increase in demand for ML skills have been Cambridge, Oxford, London, Reading and Bristol, while the industries with the greatest demand have been Financial Services, Retail, Legal & Accounting, Management Consultancy and Office Administration & Business Support.

I exploit differences in exposure of firms, local labour markets and industries to advances in ML based upon their hiring composition by occupation in 2012&13 and the extent to which advances in AI have been relevant to the tasks that different occupations conduct, as measured by Felten et al. (2019) using the EFF's AI Progress Metrics, which track progress in specific AI capabilities, such as image recognition or machine translation. I show that firms, local labour markets and industries that were more exposed to advances in ML in 2012&13 subsequently saw a relative increase in the posting of job adverts demanding ML skills between 2012 and 2017.

Using baseline ML exposure as an IV for the growth in ML demand, I then explore the impact of ML demand, as measured by the posting of online job adverts demanding ML skills, on UK labour markets and services offshoring. For UK labour markets, I find that growth in demand for ML skills has had a positive impact on the growth in demand for lower skilled occupations within the same firm, but has had no effect on the growth in demand for higher skilled occupations. There is some evidence that this former effect dominates and the net impact on total non-ML vacancies within firms is positive, particularly when comparing firms within the same industry. Growth in demand for ML skills has had a negative effect on the demand for the occupations that are the most potentially replaceable by ML, suggesting that substitution does occur between occupations that can be automated using ML and ML expertise. However, there is also some evidence that the growth in demand for ML skills has had a larger, positive impact on the posting of vacancies for non-ML exposed occupations.

I combine these measures of the demand for ML skills with a relatively unique dataset of UK services imports to evaluate how ML deployment has affected services offshoring from the UK to LMI countries. I find that growth in ML demand has actually caused an increase in services offshoring, with particular evidence of an increase in offshoring to LMI countries, and particularly India, which has experienced a more than doubling of the value of services offshoring from the UK between 2012 and 2017. There is less evidence of a positive impact on offshoring to HI countries. This increase in offshoring to LMI countries was concentrated in the offshoring of business services and ICT-related services, such as computer, information and telecoms services. I conclude that although there is some evidence of AI replacing automatable occupations in the UK, it also appears to generate new demand for lower skilled or non-automatable occupations and for services purchased from abroad, particularly from less developed countries. In this paper I haven't been able to delve into the detailed mechanisms at play or disentangle many of the complex effects occurring. For offshoring to developing countries, it is possible that either the productivity enhancing effect of AI reducing the cost of inputs could be at play, or AI could be enabling new innovation or AI could even be facilitating offshoring and the fragmentation of services value chains. Further

work with more granular data should explore in detail which of these mechanisms is driving this aggregate positive result for services offshoring.

Acknowledgements I am grateful to Bledi Taska and Layla O'Kane at Burning Glass Technologies for their help and support facilitating access to the job postings data. I am also grateful to Rob Seamans for sharing the data from Felten et al. (2019). I acknowledge funding from the CSAE to conduct exploratory fieldwork for this project. I am grateful for insightful comments and feedback from Abi Adams, Nick Bloom, Stefan Dercon, Beata Javorcik, Matthias Qian, Simon Quinn, Daniel Trefler, Michael Webb, Chris Woodruff and seminar participants at Oxford's CSAE Workshop and Oxford's Firms, Labour and Organisations group.

Appendices of Chapter 3

3.A Additional Definitions

Included service categories

The following ITIS service products are included in the analysis, which I have grouped into the four categories outlined below:

Business and professional services: Accountancy, auditing, bookkeeping and tax consulting services, Advertising, market research and public opinion polling services, Business management and management consulting services, Public relations services, Recuritment services, Legal services, Procurement services, Property management services, Other business and professional services, Provision of RD services, Provision of product development and testing activities, Publishing services, Financial services, Life insurance claims, Life insurance premiums, Freight insurance claims, Freight insurance premiums, Other direct insurance claims, Other direct insurance premiums, Auxilliary services, Pension service receipts, Pension service charges, Standardised guarantee service claims, Standardised guarantee service premiums.

ICT-related services: Telecommunication services, Computer services, Information services.

Technical and scientific services: Health services, Architectural services, Engineering services, Scientific and other technical services (including surveying).

Other related party: Transactions between related businesses not included elsewhere

Conclusion

This thesis has presented evidence on two key questions relating to recent changes affecting international trade and global supply chains. First, how do technological advances, primarily those in robotics and Artificial Intelligence (AI), affect trade and multinational activity, and particularly, trade with developing countries? The first chapter showed that the use of robots in firms in the manufacturing sector in Spain between 1990 and 2016 had a positive impact on imports and multinational activity involving developing countries. This effect was driven by automation enabling firms to expand production, increase productivity and a subset of firms starting to import from, or open affiliates in, less developed countries as a consequence of automation. The third chapter showed that the deployment of ML in the UK, as proxied for by the posting of job adverts demanding ML skills and expertise, had a positive impact on the offshoring of services, particularly to low and middle income countries rather than high income countries, with the main beneficiary country being India. While ML deployment reduced labour demand in the UK for occupations that conduct a high share of tasks that are replaceable using ML, it increased labour demand for lower skilled occupations and occupations that are not easily replaceable, with some evidence that the latter effect dominated the former.

There remain a number of unanswered questions and avenues for further research on this topic. The dataset used in Chapter 1 did not include information on the products offshored, whether automation was conducted in the offshore destinations or detailed information on imports of affiliates by country. This chapter therefore hasn't been able to say anything about product composition or the countries affected. There are many possibilities that could explain our results but not necessarily imply positive consequences for less developed countries hoping

to grow their manufacturing sectors. One possibility is that automation has been occurring in Spain's existing offshore production destinations. Another is that the increase in imports could have been in less labour-intensive goods, or concentrated in middle or upper middle income countries, such as China, rather than new offshoring destinations. This is in fact what I find in a separate paper, Bergmann and Stapleton (2020), using Danish data; automation increased offshoring to both low and middle income and high income countries, but for low and middle income countries, the increase was only directed at existing offshore destinations, whereas for high income countries firms sought out new countries to offshore to.

Additionally, this dataset also lacked information on the number of robots used, so we were only able to speak to the intensive margin of automation in a limited way. Hallward-Driemeier and Nayyar (2019) show that the impact of robot adoption is non-linear in the quantity of robots adopted. Further work, ideally with more detailed data on some of these areas, could explore many of these remaining questions. Finally, surrounding the micro-level implications of robots on firm outcomes, there remain many unknowns. We documented the importance of firm heterogeneity but only scratched the surface in terms of outcomes within Spain. Further work when more detailed data becomes available should explore firm dynamics and the impact of automation on competitor firms, market stealing effects, market concentration, the rise of superstar firms, profits and product quality.

The third chapter also offered early evidence on an emerging, complex and to date under-researched topic. There remains much more work that could and should follow. In a subject area with scarce data on whether, to what extent and for what purposes ML is being deployed commercially, we offered a novel way to measure the deployment of ML using job adverts. This method could be used to expand our understanding on the impacts of AI in myriad dimensions. More detailed work using text analysis could also attempt to induce the uses of AI from the job adverts. On the specific topic of services offshoring, further work could aim to go into more depth on the mechanisms through which ML deployment had a positive effect on services offshoring. Was it through the enabling effect of ML on facilitating offshoring? Or the impact of ML on firm growth and productivity, as in Chapter 1? Or demand for new tasks? Additional work could also attempt to ex-

plore this relationship at the firm level and could seek new identification strategies to evaluate the causal impact of AI.

As with Chapter 1, in this paper I was not able to explore whether ML was being deployed in offshore destinations like India. Further work could explore whether this is the case and whether there is path dependence in offshoring relationships such that firms that were initially conducting services activities offshore maintained their contracts but subsequently upgraded technology when it became available. While the result of this paper was generally positive for low and middle income countries, there remain a number of possibilities that could still mean that this result has negative implications for certain countries or groups of workers. The impact on offshoring could reflect a similar pattern to that observed in terms of labour demand in the UK in that there is a decrease in demand for replaceable occupations but an even greater increase in demand for other occupations, resulting in a net positive effect. This could imply negative outcomes for some skill groups. Further work could aim to undercover these distributional consequences within India. Additionally, this result could still imply that the advantage of labour costs has declined but that initial offshoring destinations were integrated enough into supply chains to maintain their position. Further work could explore the impact on the emergence of new offshoring destinations.

The second question addressed in this thesis was how the threat of reversing trade integration affects labour demand, as measured by the posting of online job adverts. Chapter 2 studied the impact of the Brexit referendum and showed that the threat of trade barriers on UK professional services exports to the EU had a substantial negative effect on the posting of online job adverts after the Brexit referendum, particularly for higher-skilled jobs. There are a number of avenues for further research on this topic. Further work could probe more deeply into the reasons for this decline in hiring: was it temporary and will bounce back after certainty emerges? Or have professional services firms permanently closed and relocated elsewhere? In this paper we also didn't delve into the types of work affected; additional research could do so. Finally, we did not attempt to look at other margins of labour market adjustment, such as separations. Further work could explore these areas and relate our findings to other labour market outcomes.

References

Acemoglu, Daron and David Autor, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in "Handbook of Labor Economics," Vol. 4, Elsevier, 2011, pp. 1043–1171.

_ and Pascual Restrepo, "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, June 2018, *108* (6), 1488–1542.

_ and _ , "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, August 2019.

_ and _ , "The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand," Working Paper 25682, National Bureau of Economic Research March 2019. Series: Working Paper Series.

_ , Claire LeLarge, and Pascual Restrepo, "Competing with Robots: Firm-Level Evidence from France," Working Paper 26738, National Bureau of Economic Research February 2020.

Aghion, C. Antonin, S. Bunel, and X. Jaravel, "What Are the Labor and Product Market Effects of Automation? New Evidence from France," 2020.

Aghion, Philippe, Benjamin F Jones, and Charles I Jones, "Artificial Intelligence and Economic Growth," Working Paper 23928, National Bureau of Economic Research October 2017. Series: Working Paper Series.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, *Prediction Machines: The Simple Economics of Artificial Intelligence*, Harvard Business Review Press, 2018.

_ , _ , and _ , "The Economics of Artificial Intelligence: An Agenda," May 2019.

Ahir, Hites, Nicholas Bloom, and Davide Furceri, "The World Uncertainty Index," *SSRN Electronic Journal*, 2018.

Amiti, M. and S.-J. Wei, "Fear of service outsourcing: is it justified?," *Economic Policy*, April 2005, *20* (42), 308–347.

Amiti, Mary, Stephen Redding, and David Weinstein, "The Impact of the 2018 Trade War on U.S. Prices and Welfare," *NBER WORKING PAPER SERIES*, March 2019, *Working Paper 25672.*

Antràs, Pol, "Conceptual Aspects of Global Value Chains," *World Bank Economic Review*, 2020.

_ and Stephen R. Yeaple, "Multinational Firms and the Structure of International Trade," in "Handbook of International Economics," Vol. 4, Elsevier, 2014, pp. 55–130.

Ariu, Andrea, Florian Mayneris, and Mathieu Parenti, "How services boost goods exports," February 2020. Library Catalog: VoxEU.

_ , Holger Breinlich, Gregory Corcos, and Giordano Mion, "The interconnections between services and goods trade at the firm-level," *Journal of International Economics*, January 2019, *116*, 173–188.

Arnold, Jens M., Beata S. Javorcik, and Aaditya Mattoo, "Does services liberalization benefit manufacturing firms?," *Journal of International Economics*, September 2011, *85* (1), 136–146.

Arnold, Jens Matthias, Beata Javorcik, Molly Lipscomb, and Aaditya Mattoo, "Services Reform and Manufacturing Performance: Evidence from India," *The Economic Journal*, February 2016, *126* (590), 1–39.

Artuc, Erhan, Luc Christiaensen, and Hernan Jorge Winkler, "Does Automation in Rich Countries Hurt Developing Ones? : Evidence from the U.S. and Mexico," Technical Report 8741, The World Bank February 2019.

_ , Paulo S. R. Bastos, and Bob Rijkers, "Robots, Tasks and Trade," Technical Report 8674, The World Bank December 2018.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms," *Quarterly Journal of Economics*, 2020, *135* (2).

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, October 2013, *103* (6), 2121–2168.

_ , _ , _ , and Jae Song, "Trade Adjustment: Worker-Level Evidence," *Quarterly Journal of Economics*, November 2014, *129* (4), 1799–1860.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, November 2016, *131* (4), 1593–1636.

Baldwin, Richard, *The Great Convergence: Information Technology and the New Globalization*, Harvard University Press, 2016.

_ , "The Globotics Upheaval: Globalisation, Robotics and the Future of Work," 2019.

_ and Rikard Forslid, "Globotics and Development: When Manufacturing is Jobless and Services are Tradable," Working Paper 26731, National Bureau of Economic Research February 2020.

Barone, Guglielmo and Federico Cingano, "Service Regulation and Growth: Evidence from OECD Countries," *The Economic Journal*, September 2011, *121* (555), 931–957.

Benzell, Seth, Guillermo LaGarda, and Daniel Rock, "Do Labor Demand Shifts Occur Within Firms or Across Them? Non-Routine-Biased Technological Change, 2000-2016," *SSRN Journal*, 2019.

Bergmann, Friedrich and Katherine Stapleton, "Technology and Global Value Chains: Evidence from Denmark," 2020.

Bernanke, Ben S., "Irreversibility, Uncertainty, and Cyclical Investment," *The Quarterly Journal of Economics*, February 1983, *Volume 98* (Issue 1), 85–106.

Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott, "The Margins of US Trade," *American Economic Review*, May 2009, *99* (2), 487–493.

_ , Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum, "Plants and Productivity in International Trade," *American Economic Review*, September 2003, *93* (4), 1268–1290.

Bernard, Andrew B, Teresa C Fort, Frederic Warzynski, and Valerie Smeets, "Heterogeneous Globalization: Offshoring and Reorganization," 2019, p. 39. Series: Working Paper Series.

Bessen, James E., Maarten Goos, Anna Salomons, and Wiljan Van den Berge, "Automatic Reaction - What Happens to Workers at Firms that Automate?," SSRN Scholarly Paper ID 3328877, Social Science Research Network, Rochester, NY January 2019.

Beverelli, Cosimo, Matteo Fiorini, and Bernard Hoekman, "Services trade policy and manufacturing productivity: The role of institutions," *Journal of International Economics*, January 2017, *104*, 166–182.

Blanchard, Emily J, Chad P Bown, and Davin Chor, "Did Trump's Trade War Impact the 2018 Election?," *NBER Working Paper*, November 2019, *26434*, 35.

Blinder, Alan S. and Alan B. Krueger, "Alternative Measures of Offshorability: A Survey Approach," *Journal of Labor Economics*, 2013, *31* (2), S97–S128.

Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, 2009, *77* (3), 623–685.

_ , Philip Bunn, Scarlet Chen, Paul Mizen, Pawel Smietanka, Greg Thwaites, and Garry Young, "Brexit and uncertainty: insights from the Decision Maker Panel," *Bank of England, Staff Working Paper No. 780*, February 2019.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-share Research Designs," 2019.

Breinlich, Holger and Chiara Criscuolo, "International trade in services: A portrait of importers and exporters," *Journal of International Economics*, July 2011, *84* (2), 188–206.

_ , Anson Soderbery, and Greg C. Wright, "From Selling Goods to Selling Services: Firm Responses to Trade Liberalization," *American Economic Journal: Economic Policy*, November 2018, *10* (4), 79–108.

_ , Elsa Leromain, Dennis Novy, and Thomas Sampson, "Voting with Their Money: Brexit and Outward Investment by UK Firms," *CESifo Working Paper no. 7751*, July 2019, p. 30.

Bresnahan, Timothy, "General Purpose Technologies," Handbook of the Economics of Innovation, Elsevier 2010.

Bresnahan, Timothy F. and M. Trajtenberg, "General purpose technologies 'Engines of growth'?," *Journal of Econometrics*, January 1995, *65* (1), 83–108.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics," Working Paper 24001, National Bureau of Economic Research November 2017. Series: Working Paper Series.

_ , Tom Mitchell, and Daniel Rock, "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?," *AEA Papers and Proceedings*, May 2018, *108*, 43–47.

_ , Xiang Hui, and Meng Liu, "Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform," *Management Science*, September 2019, *65* (12), 5449–5460. Publisher: INFORMS.

Bustos, Paula, "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *American Economic Review*, February 2011, *101* (1), 304–340.

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang, "Tariff Passthrough at the Border and at the Store: Evidence from US Trade Policy," *Working Paper*, May 2019, p. 16.

Chen, Cheng and Claudia Steinwender, "Import Competition, Heterogeneous Preferences of Managers, and Productivity," Working Paper 25539, National Bureau of Economic Research February 2019. Series: Working Paper Series.

Cheng, Hong, Ruixue Jia, Dandan Li, and Hongbin Li, "The Rise of Robots in China," *Journal of Economic Perspectives*, May 2019, *33* (2), 71–88.

Chui, Michael and Sankalp Malhotra, "Adoption of AI advances, but foundational barriers remain | McKinsey," 2019. Library Catalog: www.mckinsey.com.

Cockburn, Iain M, Rebecca Henderson, and Scott Stern, "The Impact of Artificial Intelligence on Innovation," Working Paper 24449, National Bureau of Economic Research March 2018. Series: Working Paper Series.

Costa, Rui, Swati Dhingra, and Stephen Machin, "Trade and Worker Deskilling," *CEP Discussion Paper Series*, 2019.

Crinò, Rosário, "Service Offshoring and White-Collar Employment," *The Review of Economic Studies*, 2010, p. 39.

Crowley, Meredith, Ning Meng, and Huasheng Song, "Tariff scares: Trade policy uncertainty and foreign market entry by Chinese firms," *Journal of International Economics*, September 2018, *114*, 96–115.

_ , Oliver Exton, and Lu Han, "Renegotiation of Trade Agreements and Firm Exporting Decisions: Evidence from the Impact of Brexit on UK Exports," 2018, p. 35.

Deming, David and Kadeem Noray, "STEM Careers and the Changing Skill Requirements of Work," *NBER Working Paper Series*, September 2018.

_ and Lisa B Kahn, "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals," *NBER Working Paper Series*, 2017, p. 33.

Dixit, Avinash K and Robert S Pindyck, *Investment under uncertainty*, Princeton, N.J.: Princeton University Press, 1994. OCLC: 777593629.

Ebell, Monique, "Assessing the Impact of Trade Agreements on Trade," *National Institute Economic Review*, November 2016, *238*, R31–R42.

Eppinger, Peter S., "Service offshoring and firm employment," *Journal of International Economics*, March 2019, *117*, 209–228.

Faber, Marius, "Robots and reshoring: Evidence from Mexican local labor markets," Technical Report 2018/27, Faculty of Business and Economics - University of Basel October 2018. Publication Title: Working papers.

Fajgelbaum, Pablo, Pinelopi Goldberg, Patrick Kennedy, and Amit Khandelwal, "The Return to Protectionism," *NBER WORKING PAPER SERIES*, October 2019, *Working Paper 25638*.

Feenstra, Robert C. and Gordon H. Hanson, "The Impact of Outsourcing and High-Technology Capital on Wages: Estimates For the United States, 1979–1990," *Q J Econ*, August 1999, *114* (3), 907–940.

Felten, Edward, Manav Raj, and Robert Seamans, "The Occupational Impact of Artificial Intelligence: Labor, Skills, and Polarization by Edward W. Felten, Manav Raj, Robert Seamans :: SSRN," 2019.

Fetzer, Thiemo, "Did Austerity Cause Brexit?," *American Economic Review*, November 2019, *109* (11), 3849–3886.

Fort, Teresa C., "Technology and Production Fragmentation: Domestic versus Foreign Sourcing," *Rev Econ Stud*, April 2017, *84* (2), 650–687.

Freund, Caroline, Alen Mulabdic, and Michele Ruta, *Is 3D Printing a Threat to Global Trade? The Trade Effects you Didn't Hear About* Policy Research Working Papers, The World Bank, September 2019.

Frey, Carl Benedikt and Michael A. Osborne, "The future of employment: How susceptible are jobs to computerisation?," *Technological Forecasting and Social Change*, January 2017, *114*, 254–280.

Gervais, Antoine and J. Bradford Jensen, "The tradability of services: Geographic concentration and trade costs," *Journal of International Economics*, May 2019, *118*, 331–350.

Goldfarb, Avi and Daniel Trefler, "AI and International Trade," Working Paper 24254, National Bureau of Economic Research January 2018. Series: Working Paper Series.

_ , Bledi Taska, and Florenta Teodoidis, "Could Machine Learning be a General-Purpose Technology? Evidence from Online Job Postings," 2020.

Graetz, Georg and Guy Michaels, "Robots at Work," *The Review of Economics and Statistics*, July 2018, *100* (5), 753–768.

Graziano, Alejandro, Kyle Handley, and Nuno Limão, "Brexit Uncertainty and Trade Disintegration," December 2018.

Hallward-Driemeier, Mary and Gaurav Nayyar, "Have Robots Grounded the Flying Geese?: Evidence from Greenfield FDI in Manufacturing," SSRN Scholarly Paper ID 3510400, Social Science Research Network, Rochester, NY December 2019.

Hallward-Driemeier, Mary C. and Gaurav Nayyar, *Trouble in the Making?: The Future of Manufacturing-Led Development* 2017.

Handley, Kyle and Nuno Limão, "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States," *American Economic Review*, September 2017, *107* (9), 2731–2783.

Harrigan, James, Ariell Reshef, and Farid Toubal, "The March of the Techies: Job Polarization Within and Between Firms," *Research Policy*, June 2020, p. 104008.

Hassler, John A.A., "Variations in risk and fluctuations in demand: A theoretical model," *Journal of Economic Dynamics and Control*, June 1996, *20* (6-7), 1115–1143.

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.

Honnibal, Matthew and Mark Johnson, "An Improved Non-monotonic Transition System for Dependency Parsing," in "Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing" Association for Computational Linguistics Lisbon, Portugal September 2015, pp. 1373–1378.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang, "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review*, June 2014, *104* (6), 1597–1629.

IFR, "World Robotics 2019 Edition," Technical Report 2019.

IMF, "Balance of Payments Manual," 1993. Library Catalog: www.imf.org.

Jensen, J Bradford and Lori Kletzer, "Tradable Services: Understanding the Scope and Impact of Services Outsourcing," *Peterson Institute for International Economics*, September 2005, *Working Paper* (05-9).

Jensen, Robert, "Do Labor Market Opportunities Affect Young Women's Work and Family Decisions? Experimental Evidence from India," *Q J Econ*, May 2012, *127* (2), 753–792. Publisher: Oxford Academic.

Jensen, Robert T, "Economic Opportunities and Gender Differences in Human Capital: Experimental Evidence for India," Working Paper 16021, National Bureau of Economic Research May 2010. Series: Working Paper Series.

Klinger, J., J. Mateos-Garcia, and K. Stathoulopoulos, "Deep learning, deep change? Mapping the development of the Artificial Intelligence General Purpose Technology," *arXiv:1808.06355 [cs, econ]*, August 2018. arXiv: 1808.06355.

Lazear, Edward P. and James R. Spletzer, "Hiring, Churn, and the Business Cycle," *American Economic Review*, May 2012, *102* (3), 575–579.

Lederman, Daniel, Marcelo Olarreaga, and Lucy Payton, *Export Promotion Agencies Revisited* Policy Research Working Papers, The World Bank, November 2009.

Li, Xin, Bo Meng, and Zhi Wang, "Recent patterns of global production and GVC participation," in "in" April 2019, pp. 19–54.

Lin, Justin, "How to Seize the 85 million Jobs Bonanza," 2011. Library Catalog: blogs.worldbank.org.

Liu, Runjuan and Daniel Trefler, "A sorted tale of globalization: White collar jobs and the rise of service offshoring," *Journal of International Economics*, May 2019, *118*, 105–122.

Mattoo, Aaditya, Nadia Rocha, and Michele Ruta, "The Evolution of Deep Trade Agreements," *World Bank Policy Research Working Paper 9283*, 2020.

Mayer, Thierry, Vincent Vicard, and Soledad Zignago, "The Cost of Non-Europe, Revisited," *CEPII Working Paper*, 2018, p. 50.

McMillan, Margaret, Dani Rodrik, and Claudia Sepulveda, "Structural Change, Fundamentals, and Growth : A Framework and Case Studies," 2017.

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

Miller, George A., "WordNet: a lexical database for English," *Commun. ACM*, November 1995, *38* (11), 39–41.

MMC Ventures and Barclays, "The State of AI Report 2019," Technical Report 2019. Library Catalog: www.stateofai2019.com.

Mulabdic, Alen, Alberto Osnago, and Michele Ruta, "Deep Integration and UK-EU Trade Relations," *World Bank Working Paper*, 2017, p. 24.

NASSCOM, "Facts & Figures," August 2018. Library Catalog: www.nasscom.in.

New Financial, "Report: Brexit & the City - the impact so far," March 2019. Library Catalog: newfinancial.org Section: Non classé.

Office for National Statistics, "International Trade in Services, 1996-2017: Secure Access. [Data Collection]," 2019, *8th Edition.* UK Data Service. SN: 6711, http://doi.org/10.5255/UKDA-SN-6711-9 .

_ , "Annual Business Survey, 2008-2018: Secure Access. [data collection]," 2020, *13th Edition.* UK Data Service. SN: 7451, http://doi.org/10.5255/UKDA-SN-7451-13 .

_ , "Business Register and Employment Survey, 2009-2018: Secure Access. [data collection]," 2020, *8th Edition.* UK Data Service. SN: 7463, http://doi.org/10.5255/UKDA-SN-7463-8.

ONS, "International trade in services, UK - Office for National Statistics," 2018.

_ , "UK Balance of Payments, The Pink Book: 2019," 2019.

Oster, Emily and Bryce Millett Steinberg, "Do IT service centers promote school enrollment? Evidence from India," *Journal of Development Economics*, September 2013, *104*, 123–135.

Oxford English Dictionary, *Oxford English Dictionary*, Oxford University Press, 2020. Library Catalog: www.oed.com.

Pedemonte, Mathieu, Tara Vishwanath, and Roman D Zarate, "Trade, Robots and Automation: The Impact of US Robots on Labor Outcomes in Developing Countries," 2019, p. 31.

Pierce, Justin R. and Peter K. Schott, "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, July 2016, *106* (7), 1632–1662.

Rock, Daniel, "Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence," *SSRN Journal*, 2019.

Seamans, Robert, "We Won't Even Know If A Robot Takes Your Job," 2017. Library Catalog: www.forbes.com Section: Business.

_ and Manav Raj, "AI, Labor, Productivity and the Need for Firm-Level Data," 2018, p. 14.

Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter, "Firming Up Inequality," *Q J Econ*, February 2019, *134* (1), 1–50.

South China Morning Post, "Rise of the robots: 60,000 workers culled from just one factory as China's struggling electronics hub turns to artificial intelligence," May 2016. Library Catalog: www.scmp.com Section: News.

Stanford AI Index, "AI Index 2019 | Stanford HAI," 2019.

Stapleton, Katherine, "Automation, Global Value Chains and Development: What do we know so far?," *Pathways for Prosperity Commission Background Paper Series; no. 26.*, 2019, p. 23.

Steinwender, Claudia, "Real Effects of Information Frictions: When the States and the Kingdom Became United," *American Economic Review*, March 2018, *108* (3), 657–696.

Tambe, Prasanna and Lorin M. Hitt, "The Productivity of Information Technology Investments: New Evidence from IT Labor Data," *Information Systems Research*, March 2012, *23* (3-part-1), 599–617. Publisher: INFORMS.

The Financial Express, "After Paul Krugman warns of 'mass unemployment', Raghuram Rajan says AI threatens skilled jobs," March 2018. Library Catalog: www.financialexpress.com.

UNCTAD, "World Investment Report 2019 - Special Economic Zones," *World Investment Report*, 2019, p. 80.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," *SSRN Journal*, 2019.

World Bank, "Philippines Economic Update 2018," Technical Report 2018.

_ , "World Development Indicators | DataBank," 2020.