The London School of Economics and Political Science

Essays on International Trade and Firm Organization

Giuseppe Berlingieri

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy, London, December 2013.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work.

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 71,450 words.

London, December 5, 2013

Giuseppe Berlingieri

Contents

Acknowledgement

I am greatly indebted to my supervisors Luis Garicano and Gianmarco Ottaviano, and my advisor Emanuel Ornelas for invaluable guidance and advice. I am also very thankful to Stephen Redding and Daniel Sturm who advised me during the first part of my PhD.

Many individuals have provided feedback on the chapters of this thesis. I would like to thank Rachel Ngai and Silvana Tenreyro for their numerous comments and suggestions on Chapter 1, Mirabelle Muûls for her help with the data of Chapter 2, and John Romalis for the support on Chapter 3. I have also benefited from feedback and comments from Swati Dhingra, Guy Michaels, John Morrow, Veronica Rappoport, Thomas Sampson, as well as colleagues and seminar participants at the Centre for Economic Performance (CEP) and the LSE.

Financial support from UniCredit & Universities Foundation through the "Crivelli" scholarship is very gratefully acknowledged.

Many friends I have met during these years have made the PhD experience less painful and much more enjoyable, especially Patrick Blanchenay, Andrea Buffa, Luca Fornaro, Fadi Hassan, Michele Piffer, Rosa Sanchis-Guarner, and Dimitri Szerman just to name few. And the fellow PhD students of the CEP Globalisation group have made the new 32L - 2.35 open space an incredibly inspirational working environment, many thanks to Johannes Boehm, Ben Faber, Jason Garred, Frank Pisch, and Claudia Steinwender.

This thesis would not have been possible without Ceci and my family.

Abstract

This thesis analyses the impact of globalisation on the boundary of the firm and, in turn, how outsourcing policies have shaped the reallocation of labour across sectors.

The first chapter ("Outsourcing and the Rise in Services") investigates the impact of outsourcing on sectoral reallocation in the U.S. over the period 1947-2007. Roughly 40% of the growth of the service sector comes from professional and business services, an industry highly specialized in the production of intermediates and where most of the service outsourcing activity is concentrated. As a result, business services have experienced an almost fourfold increase in their forward linkage, the largest change among all industries. I find that the overall change in input-output structure of the economy accounts for 33% of the increase in service employment, and business services outsourcing contributes almost half of that amount.

The second chapter ("Exporting, Coordination Complexity, and Service Outsourcing") investigates the determinants of service outsourcing, and professional and business services in particular. Drawing on the insights of a model of the boundary of the firm based on adaptation costs and diminishing return to management, I argue that an increase in coordination complexity (e.g.: more inputs in the production process) leads firms to outsource a higher share of their total costs and to focus on their core competences. Since country-specific inputs are needed to export to a particular country (e.g.: a specific advertisement campaign), I proxy coordination complexity with the number of export destination markets and I find support for the theory using an extensive dataset of French firms. Over time, firms that export to more countries increase the amount of purchased business services; the finding is very strong and robust to size and many other determinants of outsourcing proposed in the literature. The firm-level evidence also contributes to opening the black box of fixed export costs and to establishing a new causal link between globalization and structural transformation exploiting plausibly exogenous demand shifters.

The third chapter ("Variety Growth, Welfare Gains and the Fall of the Iron Curtain") analyses two key issues in the literature of international trade: the welfare gains from trade and the estimation of the elasticity of substitution across goods. In particular I investigate the welfare gains coming from the increase in the number of varieties in the U.K. I find that the fall of the Iron Curtain and the expansion of trade with the countries of the former Soviet contribute for roughly 10% of the total gains. China, in comparison, accounts for 5% of the gains. The methodology is an improved version of the one proposed by Broda and Weinstein (2006) and Feenstra (1994), which is more robust to the definition of goods and to the classification used.

Chapter 1

Outsourcing and the Rise in Services

1.1 Introduction

The process of economic development is characterized by the reallocation of resources across the broad sectors of agriculture, manufacturing and services. As Kuznets noted in his Nobel Prize lecture, restricting attention to advanced stages of development, structural transformation coincides with the rise of the service sector and the decline of manufacturing.[1] In the U.S., the service sector (including government) today accounts for more than 83% of total employment, compared to 60% in 1947. In order to explain structural change in recent years, it is therefore key to understand the reasons behind the remarkable rise in services.

The literature on structural transformation has mainly focused on final demand channels. Yet final demand is not the only driver of the increase in services, as firms are in turn 'consumers' of goods and services through intermediate inputs. A closer look at the data reveals that a large share of the growth of the service sector is explained by industries for which final demand plays a relatively small role, namely professional and business services, finance and real estate.[2] In particular, professional and business services account for roughly 40% of the total growth, both in terms of total GDP and total employment; when finance and real estate are added, these three industries account for 50% of the service sector growth in terms of employment and 94% in terms of GDP. Starting from this basic fact, this paper analyzes the production side of the economy and the role played by firms in shaping the reallocation of labor across sectors. I propose two unexplored channels that help explain the recent rise in services: changes in the composition of intermediates and their sourcing mode.

Intermediate goods account for roughly 50% of total gross output across a large number of countries (Jones, 2011b). However, a large intermediate multiplier is not sufficient per se to affect sectoral reallocation over time: some additional variation is needed. In this paper, I first provide novel evidence for the evolution of the input-output structure of the U.S. economy over the past

[1] "The rate of structural transformation of the economy is high. Major aspects of structural change include the shift away from agriculture to non-agricultural pursuits and, recently, away from industry to services." Lecture to the memory of Alfred Nobel, December 11, 1971.

[2] Professional and business services include accounting, engineering, consulting, legal services but also maintenance, janitorial services just to cite few. See the Appendix for the precise definition. I will also refer to this industry as business services or PBS.

60 years. In particular, I show that the most important changes are related to service sectors that are intensive in the production of intermediates. For instance professional and business services have experienced an almost fourfold increase in their forward linkage, a measure of the interconnection of an industry to the rest of the economy through the supply of intermediate inputs. Second, by providing a simple gross output accounting model that captures the full sectoral linkages of the economy, I show that changes in intermediate demand account for a significant share of the total reallocation of labor across sectors, improving the predictive power of a traditional value added model. Third, I quantify the contribution of service outsourcing to the rise of the service sector.

The strong empirical regularities unveiled by Kuznets have spurred a large body of literature, which can be divided into two main categories, depending on the explanation put forward to rationalize sectoral reallocation. The first explanation, often referred to as "utility-based" or "demand-based", highlights the role of different income elasticities for different goods and dates back to Engel.[3] The second explanation, labeled "technological" or "supply-side" and first proposed by Baumol (1967), rationalizes structural change drawing on the different rates of sectoral productivity growth and on standard homothetic preferences with a less than unitary elasticity of substitution between goods.[4]

Despite the extensive work on the subject, there is still no consensus on the empirical identification of the key economic forces that drive structural transformation, as argued by Herrendorf et al. (2013b). They show that the choice of consumer preferences is just an empirical issue and depends on how final consumption is measured. This is a key point of disagreement between the two streams of existing literature, as both mechanisms ultimately depend on the form of consumer preferences. Moreover Buera and Kaboski (2009) argue that the standard theories of structural change cannot account for the steep decline in manufacturing and rise in services in recent years, and for the large deviations between value-added shares and labor shares. This paper departs from the existing literature by analyzing the production side of the economy and proposing new channels that shape structural transformation and at the same time are unrelated to final demand.

Changes in the composition of intermediates are reflected in the structure of input-output tables. Despite the growing use of input-output data, there is no systematic evidence for the evolution of the structure of sectoral linkages over time.[5] Jones (2011b) compares the input-output structure of the U.S., Japan and China in 2000, and argues that they are not very

[3]This strand of the literature employs non-homothetic preferences to achieve non-unitary income elasticities. A non-exhaustive list of works in this area includes: Matsuyama (1992), Laitner (2000), Gollin et al. (2002), Caselli and Coleman II (2001), Restuccia et al. (2008) for two-sector models focusing on the movement of labor away from agriculture; Echevarria (1997) and Kongsamut et al. (2001) for three-sector models, where the latter authors propose a model that features both structural change and constant aggregate growth. Foellmi and Zweimüller (2008) also combine the Kaldor and Kuznets' facts in a model with hierarchic preferences.

[4]Two recent contributions that combine structural change and aggregate balanced growth are: Ngai and Pissarides (2007) in a standard three-sector model; and Acemoglu and Guerrieri (2008) in a two-sector model of high versus low capital intensive industries.

[5]Caliendo and Parro (2012), di Giovanni and Levchenko (2010), Johnson and Noguera (2012), and Jones (2011a,b) are some examples of recent works that use input-output data, but all for a given year. Acemoglu et al. (2012) look at the U.S. input-output tables for the benchmark years between 1972 and 2002 but focus on the empirical densities of the total intermediate input shares.

different: they all display a sparse pattern with a strong diagonal (output of an industry used as intermediate input in the same industry) and similar intermediate multipliers. The main difference Jones points out is that business activities are less important in China, in that they are not as widely used as in Japan and in the U.S. I find the same difference for the U.S. over time. I show that the largest change in the structure of the input-output tables involves an increase in the use of services specializing in the production intermediates, especially by manufacturing industries. Professional and business services have experienced an almost fourfold increase in their forward linkage and the use of finance and real estate has also risen, albeit to a lesser extent, with their forward linkages increasing by 83% and 42%, respectively.

I study the changes in intermediate demand in a standard growth accounting model with intermediate inputs as in Hulten (1978), expanded to capture the fully fledged input-output structure of the economy similar to Horvath (1998, 2000).[6] In this setting, not only do sectoral labor shares depend on consumption shares as in a standard value added model, but also on the input-output structure of the economy through the Leontief inverse matrix. Changes in intermediate demand therefore induce a reallocation of labor across sectors. I find that, when final demand is kept constant over time, the sole evolution of the input-output structure of the economy accounts for 33% of the total increase in service employment. Although demand-side factors are certainly important, this exercise quantifies the proposed channels in a neat and simple setting, which avoids confounding the results with the choice of data and parameters not specifically related to the forces under study. Then I allow final demand to evolve over time and show that the results are not wiped out by other channels previously discussed in the literature. In fact, accounting for intermediates improves a traditional value added model prediction for the share of services by 4.7 percentage points of total employment, an amount that corresponds to 21% of the actual increase in services over the period.

What drives the changes in the use of intermediates over time? I show that one of the key forces is outsourcing. The intuition is simple: if firms contract out part of their production processes, they will have to buy these inputs from external providers, and this change will be reflected in the data as an increase in the use of intermediates. In particular, if a manufacturing firm outsources part of its headquarter services, the intermediate use of services will increase because it is likely that these inputs will be purchased from firms specializing in services. The idea that outsourcing might drive structural transformation goes back to Fuchs (1968) but, to the best of my knowledge, it has never been formally tested in a model of structural transformation.[7] Herrendorf et al. (2013b) briefly discuss this idea, arguing that outsourcing is unlikely to play a major role. Although outsourcing alone certainly cannot explain the entire process of structural transformation, at the same time the data reveal that its impact can be sizable. In fact more than 90% of the output of professional and business services is used by firms, either as intermediate input or in the form of investment. Hence final demand plays essentially no role in the growth

[6] Recent examples that employ a framework with intermediate inputs and full sectoral linkages include, among others, Ngai and Samaniego (2009) and Caliendo and Parro (2012).

[7] Fuchs points out that: "As an economy grows, there is some tendency for specialized firms to be organized to provide the business and professional services that were formerly taken care of within manufacturing and other goods-producing firms or were neglected."

of an industry that accounts for almost half of the total rise of the service sector.

Given the high share of intermediate production and the high substitutability that characterize business services, it is common in the literature to identify the rise of this industry as an increase in outsourcing. I take a similar approach in this paper and improve on the literature by controlling for internal production. In principle input-output data do not clearly distinguish the boundary of the firm. However, in the case of business services, most of the internal production is classified in auxiliary units (headquarters), which can be excluded. I show that the increase in the demand of business services comes from transactions across the boundary of the firm, and is not matched by a parallel increase in internal production. I then quantify how much of the change in intermediate use is due to business services purchased by other firms, thereby providing an estimate of the contribution of service outsourcing to the change of sectoral employment shares.[8] I do this performing a simple counter-factual exercise that fixes the demand of business services to their 1947 level and keeps it constant over time. I find that, had firms produced all their business services in-house, the service sector employment share would have been 3 percentage points smaller, which is equivalent to 14% of the total increase in the share of services.

There is much evidence that many other types of services have been outsourced over the same period, especially bearing in mind the very long time frame of the analysis. By focusing on business services only, I therefore take a conservative approach and provide a lower bound for the contribution of outsourcing to structural change. Yet I capture a large share of the total actual contribution. For instance I find that finance, despite having experienced an almost double increase in its forward linkage and having contributed to the recent rise of macroeconomic volatility as showed by Carvalho and Gabaix (2013), does not play a major role in the reallocation of labor across sectors. A potential concern is that final demand might drive the rise in business services indirectly, with firms increasing their use of services as a result of a shift in consumers' tastes. Yet an analysis of occupational data shows that, to a first approximation, the overall composition of business services has not changed over time, supporting the view of an organizational change with a reallocation of activities across the boundaries of the firms; and even where specific activities have increased their importance over time, final demand is unlikely to play a role in that change.

The paper is organized as follows. The next section discusses the main stylized facts on the rise of the service sector and critically assesses the measure of outsourcing used in the analysis. I then outline the accounting framework in Section 1.3, and present the main results of the paper in the following section. Section 1.5 shows that the results are robust to the inclusion of traditional final demand channels. Finally section 1.6 discusses potential determinants of outsourcing and Section 1.7 concludes. The details on the data and extra results are presented in the Appendix.

[8] The definition of outsourcing is standard; in Helpman's (2006) words: "outsourcing means the acquisition of an intermediate input or service from an unaffiliated supplier". I focus on domestic outsourcing, rather than international outsourcing or offshoring.

1.2 The Rise of the Service Sector in the U.S.

Over the past 60 years, structural transformation in developed countries has mostly coincided with the impressive rise in the share of services. For instance, in the U.S., the share of services in total GDP has risen to 80% in 2007 from 60% in 1947, as displayed in Figure 1.1a (left-hand side axis). This is a well-known fact but what has not been sufficiently appreciated in the literature is that this growth is almost entirely explained by three industries only, namely Professional and Business Services (hereafter PBS), Finance and Real Estate.[9] Figure 1.1a also shows the total growth of the service sector and its components (right-hand side axis); PBS, Finance and Real Estate account for a growth of 18.8 percentage points of GDP, versus a total growth of 20.1 points. Adding Health Care, these four industries account for more than the total growth, meaning that other service sectors have seen their shares decreasing. PBS have increased their share in total GDP by 8.8 percentage points, accounting for 43.6% of the total growth of the entire service sector, the biggest contribution among all industries. The same graph drawn for employment is revealing (Figure 1.1b). PBS have grown by 9.2 percentage points of total employment, roughly the same amount in terms of GDP. On the other hand, Finance and Real Estate combined have increased their share in total employment by only 2.3 percentage points, versus a combined increase of 10.1 in terms of GDP. This highlights the asymmetric contribution of these industries; Finance and Real Estate contributed a lot in terms of value added but not that much in terms of employment. Given the importance of PBS, the rest of this section will investigate the implications of their rise on the structure of the economy and the determinants of the rise itself, which can be ascribed mainly to outsourcing.

1.2.1 The Change in the Input-Output Structure of the U.S. Economy

The PBS industry is unusual. In fact, in 2002 roughly 83% of its output was sold to firms as intermediate inputs compared to 44% for the economy as a whole; an additional 8% of its output was used for investment, while final consumption accounted for just 7%. One of the implications of these characteristics is that the remarkable growth in the share of PBS is reflected in a parallel change of the Input-Output (I-O, hereafter) structure of the economy; a fact that has been overlooked in the literature despite the widespread use of I-O data. Jones (2011b) asks the question how much the I-O structure of an economy differs across countries; his answer is "not much". Looking at the I-O matrices for the U.S., Japan and China in 2000, he notices that they all display a sparse pattern with a strong diagonal and just a few inputs that are widely used by all other sectors. The main difference is that business activities are less important in China: they are not as widely used as in Japan and in the U.S. A very similar picture holds true for the U.S. over time. Figure 1.2 shows the evolution of the total requirements table from

[9]Many authors have discussed the important contribution of PBS to job growth; see for instance Abramovsky and Griffith (2006), Abramovsky et al. (2004) for the U.K.; Goodman and Steadman (2002), and Yuskavage et al. (2006) for the U.S. But, to the best of my knowledge, no previous work has attempted to quantify the impact of PBS outsourcing on structural transformation.

1947 to 2002.[10] The main change is the significant increase in the use of PBS (sector 73) in the production of all other goods, and to a smaller extent the increase in the use of Finance (sector 70) and Real Estate (sector 71), other two industries for which final demand plays a relatively small role. The horizontal line corresponding to PBS was almost absent in 1947 but becomes more and more visible over time. This change is clearly depicted in Figure 1.3a that shows, for all commodities in the economy, the increase in the share of PBS in the total requirements.

The horizontal sum of the coefficients in the total requirements table is usually referred to as forward linkage, a measure of the interconnection of a sector to all other sectors through the supply of intermediate inputs. In light of the insights provided by Acemoglu et al. (2012), the sharp rise of the PBS forward linkage implies that this sector has greatly increased its influence on the rest of the economy and any shock to it will now propagate directly to a large part of the economy. Figure 1.3b shows, for some selected industries, the evolution of the forward linkage divided by the total number of sectors; in Acemoglu et al.'s (2012) setting, this quantity essentially corresponds to the elements of what they define "influence vector" (up to the labor share). The figure confirms that PBS have experienced a sharp increase in their forward linkage, overcoming sectors with a traditionally high forward linkage like transportation. PBS have in fact become the sector with the highest influence on the rest of the economy, considerably higher than the influence of the average or median sector. The forward linkage of the finance sector (sector 70) has also increased, although more moderately compared to PBS. This fact is in line with the results of Carvalho and Gabaix (2013), who show that the recent rise of macroeconomic volatility is largely explained by the rise of finance, or more specifically of its Domar weight. Their results are suggestive for the impact that the PBS sector might have on aggregate outcomes. Finally, this change is not a specific characteristic of the U.S economy; in fact, in Appendix 1.D, I show that the same pattern holds true for most OECD countries.[11]

The PBS intrinsic nature of being mainly specialized in the production of intermediate inputs calls for an investigation of the role of firms in driving the rise of the PBS share in total employment. In particular, changes in intermediate demand or managerial decisions like producing in-house or outsourcing affect the share of services in total intermediates, increasing the use of PBS. These channels remain unexplored in the literature of structural change, given the focus on final demand. PBS are the industry where most of the service outsourcing takes place; it is very common in the literature to identify the rise in use of PBS as an increase in outsourcing, and the same approach is taken here. There could be other explanations though: an overall increase in service activity both inside and outside the firm or, more simply, problems in precisely identifying the boundary of the firm in the data. The next sub-section provides evidence showing that the rise in the use of PBS is mainly driven by outsourcing.

[10] The total requirement table shows for each commodity at the bottom of the table the inputs required, both directly and indirectly, from all industries in the economy to produce a dollar of output. The strong diagonal in this case is obtained by construction.

[11] In the case of the U.K., Oulton (2001) reports a sharp increase of the Domar weight (the ratio of sectoral gross output to aggregate value added) for the combined sector finance and business services over the period 1979-1995.

1.2.2 The Rise in PBS and Outsourcing

The identification of the rise in PBS use with a rise in outsourcing is quite common in the literature[12], but this assumption could raise some concern since the I-O data do not clearly distinguish the boundary of the firm. The data are collected at the establishment level; hence, all the in-house services provided by the headquarters or by separate service-providing units will be accounted within services, and the increase of PBS use could just be an increase in the use of services produced by the same firm and not purchased from the market. Yet, a deeper analysis of industry data shows that most of the transactions take place across the boundaries of the firms, and they are not matched by a parallel increase of services produced inside the firms. Mainly using occupational data, Section 1.6 will provide further insights and evidence on the potential mechanisms that drive the rise in outsourcing.

Industry data, on which I-O data are based, offer two main arguments in support of the idea that the increase in PBS mostly coincides with an increase in service outsourcing. First of all, it is true that the data are collected at the establishment level, but service reporting units are classified within services only under the new NAICS classification, which was adopted in 1997. This means that for all previous years, under the SIC classification, the establishments providing support services were classified on the basis of the industry of the establishment they were serving, and not their primary activity.[13] Hence, all the establishments providing support services to manufacturing firms were classified within manufacturing, and the increase of PBS use by these firms necessarily coincided with transactions outside the boundary of the firm. Secondly, the share of value added or employment accounted by auxiliary units is remarkably constant over time, and it cannot explain the increase in the share of PBS. Figure 1.4 and Table 1.1 show the share of PBS in GDP and in total employment over time, according to the two different classifications. It is evident that their difference does not vary much over time. In fact, when the sub-sector corresponding to auxiliary establishments is removed from the NAICS data, the series look extremely similar under the two different classifications.

One could think of the creation of auxiliary units as a temporary phase in the life-cycle of a manufacturing company. At en early stage, services are performed internally. For instance the accounting, billing, and marketing activities are performed at the back of the production site; no separate unit exists and no separate records are kept, hence the production of these services does not show up in the data. Even when the company becomes bigger and sets up separate accounting and marketing departments, the production of these services will remain undetected unless separate records are kept. These services will appear in the data only at a further stage, when the company has grown further and has become a large multi-establishment enterprise, establishing a separate auxiliary unit that can charge intra-company users and even sell services to other enterprises. It is at this stage that the two classifications differ. Under

[12] Among others, see Abraham and Taylor (1996), Fixler and Siegel (1999), ten Raa and Wolff (2001) and Abramovsky and Griffith (2006).

[13] These establishments were called auxiliaries units in the SIC nomenclature. For further details see the U.S. Census Bureau Clarification Memorandum (http://1.usa.gov/104BWsf) and Office of Management and Budget (1987).

Table 1.1: Professional and Business Services (PBS)

(a) Share of GDP

	1948	1957	1967	1977	1987	1997	2007
PBS (NAICS)	3.34	4.10	4.87	5.62	8.12	10.11	12.09
Auxiliary Units	1.49	1.52	1.47	1.42	1.62	1.49	1.83
PBS - Aux.	1.86	2.58	3.40	4.19	6.49	8.62	10.26
PBS (SIC)	1.75	2.55	3.46	4.37	6.86	8.93	n.a.

(b) Share of Total Employment

	1948	1957	1967	1977	1987	1997	2007
PBS (NAICS)	3.34	3.88	4.82	6.13	9.18	12.01	12.56
Auxiliary Units	1.13	1.12	1.09	1.10	1.15	1.27	1.29
PBS - Aux.	2.21	2.76	3.73	5.03	8.03	10.74	11.26
PBS (SIC)	1.49	2.05	3.09	4.40	7.30	9.92	n.a.

Note: see notes in Figure 1.4.

SIC, this new auxiliary unit is classified according to the industry of the establishment it serves, that is manufacturing. Under NAICS, instead, the unit is classified on the basis of its primary activity, which is PBS; hence the data will display an increase of PBS intermediate use, despite coming from within the boundary of the firm. Eventually, increased economic specialization may lead the enterprise to outsource its service inputs to external providers. At this final stage both classifications will allocate these activities to PBS, and the services bought by the manufacturing enterprise will be correctly accounted as an increase of PBS intermediate use outside the boundary of the firm.

By excluding auxiliary units in the main results of the paper, I will be able to control for internal production under both classifications and hence correctly identifying outsourcing. Unfortunately the same sharp conclusion cannot be drawn for I-O data in recent years, or at least not entirely. In fact, although the Bureau of Economic Analysis (BEA) constructs I-O tables using the same definition of industries, it applies some modifications in the case of commodities.[14] As for industry data, the BEA classifies establishments according to their primary activity; occasionally, however, it identifies some secondary products and re-classifies them into other commodities, in contrast with the Economic Census that classifies everything in the industry of the primary product. This re-classification only affects small single establishment firms with one single secondary product (but large enough to keep separate records).[15] In fact, whenever two or more support activities cross six-digit NAICS industries, they are treated as auxiliary units and classified in NAICS sector 55 (Management of Companies and Enterprises),

[14] The definition of industries corresponds to the SIC or NAICS definition when the standard tables (before industry redefinitions) are used. See the Online Appendix 1.B.2.

[15] An example is a small newspaper publisher that produces advertising as its single secondary product. For further details see Horowitz and Planting (2006).

which I will exclude. This is the case for medium and large multi-establishment enterprises that usually internally produce more than one support activity.

The problem of internal transactions therefore only remains for those small firms whose secondary products are re-classified by the BEA from manufacturing to PBS. These transactions are small in absolute terms and they are unlikely to drive the results. This statement is consistent with the evidence for goods provided by Atalay et al. (2012) for the domestic operations of U.S. multi-plants firms, and by Ramondo et al. (2012) for intra-firm trade of U.S. multinational firms. Both papers show that shipments between establishments owned by the same firm are surprisingly low and extremely skewed towards towards large plants: the internal shipments of the median plant are zero or very low in both studies. Hence, by controlling for the internal transactions of medium and large plants, I am likely to capture the vast majority of internal service production recorded in the data.[16]

Moreover, there are two extra reasons to believe that the results will provide a robust estimate for outsourcing. First, I only consider PBS outsourcing, while there is much evidence that many other types of services have been outsourced, especially bearing in mind the long time frame of the analysis: transportation and warehousing are good examples.[17] Even though a small fraction of the change in PBS use accounted as outsourcing might come from internal transactions, many other types of services are not included, possibly causing an even larger bias in the opposite direction. I do not include them in the baseline results to be more conservative. In fact other services like transportation and wholesale trade are not classified within auxiliary units, hence contrary to PBS I would not be able to properly control for internal transactions. The second reason is that only the difference in service outsourcing will matter in the analysis. If the internal production of secondary products stays constant in relative terms over time, these internal transactions cannot possibly drive the result. The constant share accounted by auxiliary units, as shown in Table 1.1 and Figure 1.4, confirms this view.

This fact also provides evidence that the increase in the intermediate use of services is not a simple progressive shift towards service activity: the increase of purchased services is not matched by an equal increase of services internally produced. Or, to put it another way, even if firms started using more services for technological or other motives, they decided to purchase them from the market rather than produce them in-house. PBS are intrinsically very substitutable: for instance, a firm always has the option to employ an accountant or an engineer in-house instead of buying accounting and engineering services from specialized firms. Whether a firm today needs more accounting inputs due to the more complex regulatory environment or the firm is simply outsourcing the very same tasks it used to produce with internal employees,

[16] Appendix 1.B.2 shows that the industry redefinitions performed by the BEA have a negligible impact on the magnitude of the results. It is reasonable to assume that the commodity re-classifications, which unfortunately cannot be observed, will have a similar small effect. Moreover any re-classification that takes place within manufacturing will not matter for the analysis; only the re-classifications from manufacturing to services, and PBS in particular, are a source of concern. The only examples provided by the BEA that fall into this category are advertising and data processing services.

[17] For instance, as reported by Alvarenga and Malmierca (2010), most companies managed the physical distribution of their own products in the '50s. Then two new companies, FedEx and DHL, together with UPS, started specializing only in that and quickly their logistical skills significantly eclipsed those of many manufacturing companies. What was done in-house in the '50s now is seen as a function best performed by external providers.

it is not of primary importance for the quantitative analysis performed in this paper. Despite the option of internal production the firm decided to purchase the input from the market, so whatever the fundamental reason behind this choice may be, what is key in order to calculate the impact of outsourcing on the reallocation of labor across sectors is to correctly identify market transactions. Of course understanding why firms are outsourcing more services today is another interesting - albeit difficult - question to answer; the main problem is that it is hard to observe what a firm produces in-house. Despite a full analysis being beyond the scope of the paper, Section 1.6 will try to shed some light on this important issue and show that, to a first approximation, the overall composition of business services has not changed over time.

Overall, the analysis of industry level data supports the view that most of the increase in the share of PBS has been driven by outsourcing. Firms, and manufacturing firms in particular, have increasingly bought services from the market instead of producing them in-house, causing a reallocation of resources across sectors. Herrendorf et al. (2013b) briefly discuss the role of outsourcing in shaping structural transformation; they claim that is not a major driving force arguing that PBS account for less than half of the increase in services and that a substantial share of PBS might reflect purchases directly made by consumers. Yet final demand accounts for just 7% of total PBS output and, according to their findings, PBS account for 41.5% of the total increase in services. Even though structural transformation cannot be entirely driven by outsourcing, at the same time the data reveal that its impact can be sizable. A back-of-the-envelope calculation using their results shows that, once the share of intermediates in Finance and Real Estate is also included, more than 53% of the total change in services comes from an increase in the use of intermediates.[18] Firms can therefore play an important role in driving structural transformation, and managerial decisions like outsourcing are likely to have a sizable impact.

1.3 A Simple Gross Output Accounting Model

I use a simple accounting framework in order to quantify the contribution of the evolution of sectoral linkages, and of outsourcing in particular, to the reallocation of employment across sectors. The framework builds on standard growth accounting with intermediate inputs, widely used in the productivity literature since Hulten (1978), and expanded to capture the fully fledged I-O structure of the economy as in the work of Horvath (1998, 2000). The main aim of this study is to perform an accounting exercise and not to explain why firms are changing their sourcing behavior over time.[19] The changes in the I-O structure of the economy are therefore taken as given and simply regarded as exogenous changes in the production function. In this respect, the approach is close in spirit to the work of Carvalho and Gabaix (2013), who take the change of

[18]This result is simply obtained by summing up the contributions of Finance & Real Estate and PBS to the total increase in services, 48.8% and 41.5% respectively, weighted by the average share of intermediates in their output, which is 39% for the former and 82% for the latter. If owner-occupied dwellings are excluded, the share of intermediates in Finance & Real Estate output increases to 61% (in 2002) and the overall contribution to 64%.

[19]In order to unveil the causes of this process, it is key to understand the main reasons why firms have started outsourcing more services over time. Section 1.6 provides some insights on this important issue, but a full response to this question is beyond the scope of this paper and it is left for future research. See Chapter 2.

the Domar weights as given. The model is in a closed economy. The main reason for this choice is that, although the importance of imported services has risen in recent years, their magnitude is still very low, accounting for just 2.7% of total PBS in 2004 as reported by Yuskavage et al. (2006). This fact is also confirmed by the results of this paper, which find that imported services play a very small role in the change of the I-O structure of the economy over time. Therefore the measure of outsourcing considered in this paper almost coincides with domestic outsourcing, given that the international dimension still plays a small role in the case of services.

1.3.1 The Economic Environment

1.3.1.1 Technology and Production

There is an arbitrary number of J sectors in the economy, even though in the baseline case I will consider just three aggregate sectors: agriculture, manufacturing and services. The production function for the good in sector j is given by:

$$Y_j = A_j L_j^{\beta_j} \left[\prod_{k=1}^{J} M_{kj}^{\gamma_{kj}} \right]^{1-\beta_j} \tag{1.1}$$

where A_j is the level of productivity, L_j is the amount of labor and $\beta_j \geq 0$ is the share of value added in sector j. M_{kj} is the amount of intermediate good from sector k used to produce the good in sector j. Note that the production function employs intermediate goods potentially from all sectors; $\gamma_{kj} \geq 0$ is the share of intermediates from sector k and such that $\sum_{k=1}^{J} \gamma_{kj} = 1$ for any sector j. There is no capital in the model, so there is no dynamics and the equilibrium is simply a sequence of static economies. Hence time subscripts are not reported unless explicitly needed.

The Cobb-Douglas formulation for the production of gross output is quite common in growth accounting.[20] It is assumed here to keep the model as standard as possible and, most importantly, because it can be very easily and intuitively calibrated in the data. On the other hand, the intuition for outsourcing in its starkest form is a pure relabeling effect, according to which the same tasks previously performed inside the firm are simply outsourced to external providers. If the new supplier is classified in a different sector, for instance a manufacturing firm that contracts out its accounting to a specialized service provider, this will bring about a reallocation of resources across sectors. Under this interpretation, the outsourced task is considered as essentially the same, regardless whether it is produced inside or outside the firm.[21] Therefore the Cobb-Douglas formulation is not the ideal one, as one would think of those tasks as almost perfectly substitutable. Nevertheless, for the reasons just outlined, the production function is

[20] Ngai and Pissarides (2007) show that a Cobb-Douglas functional form is needed in order to obtain a balanced growth path. Herrendorf et al. (2013a) find that a Cobb-Douglas production function well captures U.S. postwar structural transformation, and even more so in a gross output framework like the present one.

[21] Notice that this very stark interpretation is not the only explanation; outsourcing can in fact take several forms. For instance outsourcing could entail the substitution of an old superseded task with a new more technologically advanced one. In this sense outsourcing could be a way of accessing new technologies that would be too costly to be produced in-house, as Bartel et al. (2012) have argued. See Section 1.6.

assumed to be Cobb-Douglas and the perfect substitutability is imposed through some simple counterfactual exercises, which are described at the end of this section.

Each sectoral good can be either consumed or used as an intermediate in the production of the other goods according to (1.1), so the market clearing for each sector requires:

$$Y_j = C_j + \sum_{k=1}^{J} M_{jk} \tag{1.2}$$

where C_j is consumption of good j. Households are endowed with L units of labor that supply inelastically at the rental price w. All factor and goods markets are characterized by perfect competition and labor is perfectly mobile across sectors. Producers of each good solve the following problem:

$$\min_{L_j,\{M_{kj}\}_{k=1}^{J}} wL_j + \sum_{k=1}^{J} P_k M_{kj} \qquad \text{s.t.} \qquad A_j L_j^{\beta_j} \left[\prod_{k=1}^{J} M_{kj}^{\gamma_{kj}} \right]^{1-\beta_j} \geq Y_j \tag{1.3}$$

The conditional factor demands are:

$$L_j = \beta_j \frac{P_j Y_j}{w} \tag{1.4}$$

$$M_{kj} = \gamma_{kj}(1-\beta_j)\frac{P_j Y_j}{P_k} \tag{1.5}$$

1.3.1.2 Sectoral Labor (Re-)Allocation

Using the good market clearing condition in (1.2) and the equilibrium demand for intermediates according to (1.5), it is possible to get an expression for the value of gross output for each sector j as follows:

$$P_j Y_j = P_j C_j + P_j \sum_{k=1}^{J} M_{jk} = P_j C_j + \sum_{k=1}^{J} \gamma_{jk}(1-\beta_k) P_k Y_k \tag{1.6}$$

Using the equilibrium demand for labor according to (1.4), the labor share l_j of each sector can be written as follows:

$$l_j = \frac{L_j}{L} = \frac{\beta_j P_j Y_j}{wL} = \beta_j X_j + \beta_j \sum_{k=1}^{J} \gamma_{jk} \left(1-\beta_k\right) \frac{P_k Y_k}{wL} \tag{1.7}$$

where $X_j = \frac{P_j C_j}{wL}$ is the consumption expenditure share of sector j, or, using a terminology more consistent with the empirical application, the final uses expenditure share.[22] Therefore the labor shares reflect the presence of intermediates and the interrelation of sectors. In fact, the labor share of sector j depends on the value of gross output of all other sectors $\{P_k Y_k\}_{k=1}^{J}$, and the overall intensity with which each other sector uses the output from sector j as intermediate input, $\gamma_{jk}\left(1-\beta_k\right)$.

[22]Only final consumption is explicitly modeled, but in the empirical implementation other final uses are considered as well, like government consumption and investment. Final uses would therefore be the appropriate terminology. Nevertheless, the two terms are used interchangeably in the rest of the paper.

Note that equation (1.6) forms a system of J equations; it is convenient to re-write and solve it using matrix algebra as follows:

$$\mathbf{Y} = \Omega^{-1}\mathbf{C} \tag{1.8}$$

where:

$$\mathbf{Y} = \begin{pmatrix} P_1Y_1 \\ \vdots \\ P_JY_J \end{pmatrix} \quad \mathbf{C} = \begin{pmatrix} P_1C_1 \\ \vdots \\ P_JC_J \end{pmatrix} \quad \Omega = \begin{pmatrix} 1-\gamma_{11}(1-\beta_1) & \cdots & -\gamma_{1J}(1-\beta_J) \\ \vdots & \ddots & \vdots \\ -\gamma_{J1}(1-\beta_1) & \cdots & 1-\gamma_{JJ}(1-\beta_J) \end{pmatrix} \tag{1.9}$$

The matrix Ω is a J by J matrix and it can be expressed as $\Omega = I - D$, where I is an identity matrix and D is an industry-by-industry direct requirement matrix with a generic element defined as $d_{j,k} = \gamma_{jk}(1-\beta_k)$. Ω^{-1} is referred to as the total requirements table, or the Leontief inverse matrix, and can be directly obtained from I-O data. Having solved for gross output, the vector of labor shares is obtained as follows:

$$\boldsymbol{l} = \frac{1}{wL}\boldsymbol{\beta}\mathbf{Y} = \boldsymbol{\beta}\frac{\Omega^{-1}\mathbf{C}}{wL} = \boldsymbol{\beta}\Omega^{-1}\mathbf{X} \tag{1.10}$$

where:

$$\boldsymbol{l} = \begin{pmatrix} l_1 \\ \vdots \\ l_J \end{pmatrix} \quad \boldsymbol{\beta} = \begin{pmatrix} \beta_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \beta_J \end{pmatrix} \quad \mathbf{X} = \begin{pmatrix} X_1 \\ \vdots \\ X_J \end{pmatrix} \tag{1.11}$$

Therefore the labor shares differ from consumption expenditure shares due to the fully fledged I-O structure of the economy, captured by the total requirement table. The labor share in each sector is, in general, a function of the consumption share of all other sectors.

Introducing time subscripts, equation (1.10) can be re-written as:

$$\boldsymbol{l}_t = \boldsymbol{\beta}_t\Omega_t^{-1}\mathbf{X}_t \tag{1.12}$$

The sectoral labor shares can evolve for two main reasons: either because of changes in final uses, $\mathbf{X}_t$, as the literature on structural transformation has highlighted so far; or because the I-O structure of the economy changes over time. Note that the latter channel can affect employment shares in isolation, even if consumption expenditure shares do not change. This is precisely what I do in the main results of the paper: I keep final uses constant and simply evaluate the impact of the evolution of the I-O structure on the sectoral labor shares taking the matrices $\boldsymbol{\beta_t}$ and Ω_t from the data.[23] Then, as a robustness check, I allow for consumption shares to vary over time and show that the main results are not wiped out by the standard channels proposed in the literature. To this purpose preferences will be introduced in Section 1.5.

[23] This is obviously not possible in a value added model: if the share of intermediates is zero in all sectors, the matrix Ω is an identity matrix and the labor shares coincide with the final uses shares: $\boldsymbol{l}_t = \mathbf{X}_t$ if $\beta_j = 1, \forall j \in J$.

1.3.2 Accounting for Outsourcing: Three Simple Counterfactual Exercises

In order to quantify the contribution of outsourcing to structural change, I perform three counterfactual exercises. The first one consists in fixing the I-O coefficients for manufacturing to their 1947 level, which implies taking the values for the elements of the direct requirement matrix $\{d_{j,m}\}_{j=1}^{J}$ in 1947 and keeping them fixed over time. This exercise shows what would have happened to sectoral employment shares, had manufacturing firms not changed their intermediate demand over time. In the data, the importance of services in the total intermediates of the manufacturing sector ($d_{s,m}$) has strongly risen over time. Therefore fixing this coefficient to its 1947 level implies a lower labor share for the service sector, as equation (1.7) shows. The difference with the predictions obtained allowing for the full change in the I-O structure can be regarded as an upper bound for the contribution of outsourcing to sectoral reallocation. It would in fact correspond to assuming the whole increase in the use of service intermediates by manufacturing firms as coming from outsourcing. Not only are PBS included, but all other possible types of services like transportation, wholesale trade, health care, government inputs, etc... Although slightly overstretched, this is not totally implausible, as outsourcing is indeed observed even outside the PBS industry; finance, transportation and warehousing are all good examples of services that have been increasingly outsourced over time.[24] The second and third exercises are very similar; instead of fixing the direct requirements coefficients from all other sectors, only the share of inputs coming from PBS and the one coming from finance are fixed, one at a time. Table 1.2 summarizes the exercises.

Table 1.2: Counterfactual Exercises

1: No Service Outsourcing	2: No PBS Outsourcing	3: No Finance Outsourcing
$d^t_{j,m} = d^{1947}_{j,m} \quad \forall j \in J$	$d^t_{PBS,m} = d^{1947}_{PBS,m}$	$d^t_{f,m} = d^{1947}_{f,m}$

Note: m=manufacturing, f=finance.

The counterfactual corresponding to fixing the share of PBS inputs is the main focus of the paper; it answers the question of what would have happened if the share of PBS intermediate inputs to manufacturing had been fixed at its 1947 level and all PBS had been produced internally within manufacturing. Of course this exercise is correct only if the rise in PBS use comes from market transactions outside the boundary of the firm, otherwise it would not be possible to identify the result of this counterfactual as the contribution of outsourcing. Due to the re-classifications performed by the BEA, it is not possible to completely rule out the eventuality that a few transactions may come from establishments within the same firm. But, as already noted in Section 1.2.2, this problem only affects those small firms whose secondary products are re-classified from manufacturing to PBS by the BEA. As for the case of redefinitions discussed in Appendix 1.B.2, it is very unlikely that these re-classifications can have a big impact on the results. Moreover, as it is clear from the case of transportation, outsourcing also takes place in other sectors within services, especially because the focus is on the total change since 1947 and it is well documented that many of other types of services were performed internally at the

[24] See footnote 17.

beginning of the period. All in all, the contribution of PBS outsourcing is a reasonable estimate and possibly a lower bound for the overall contribution of service outsourcing.

1.4 Sectoral Reallocation in the U.S., 1948-2002

I use the accounting model outlined in the previous section to predict structural transformation in the U.S. The advantage of using U.S. data is the very long time span; I-O tables are in fact available dating from 1947. Hence, compared to other countries, it is possible to investigate sectoral reallocation over a time horizon that is long enough to display the clear pattern of structural change. This section shows that it is possible to keep the final uses expenditure shares constant and still get a positive sectoral reallocation, by allowing the I-O structure of the economy to change over time. By shutting down the final demand channel, the only driving forces come from the production side. This setting is therefore a neat environment in which to investigate the role played by firms in shaping the reallocation of labor across sectors, and in particular quantify the contribution of changes in the composition of intermediates and their sourcing mode.

1.4.1 Calibration

Following most of the literature on structural transformation, I consider three sectors in the baseline case: agriculture, manufacturing and services; hence $J = 3$ and $j \in \{a, m, s\}$. This choice implies that all the total requirements tables have to be aggregated up to three sectors only.[25] I calibrate final uses to match the employment shares in 1948, the first year for which employment data are available. Inverting equation (1.12) it is possible to get the final uses shares from the employment shares according to:

$$\mathbf{X}_t = \Omega_t \boldsymbol{\beta}_t^{-1} \boldsymbol{l}_t \tag{1.13}$$

This is the only step required to predict the evolution of labor shares when the contribution of outsourcing and of the evolution of sectoral linkages are analyzed in isolation. In fact, by keeping final uses shares constant over time, I only need data on Ω_t and $\boldsymbol{\beta}_t$ to predict labor shares according to (1.12).[26]

Armed with data from the I-O tables, I then predict employment shares until 2002. In recent years, the I-O tables are available annually, not just for the benchmark years; hence the analysis can be extended until 2007 and not just until 2002, the last benchmark year. However, some caution in interpreting the results is needed. In fact, the annual tables are computed using more aggregate data and they do not match the statistical quality of tables in benchmark years. In particular, the intermediate inputs at the detail level are estimated assuming the industry

[25] I consider a more disaggregated level only for the counterfactual exercises, in order to account for the specific PBS and finance shares.

[26] These matrices are directly available for all benchmark years, while I use interpolated values for all other years. Further details on the data and on the methodology are contained in the Appendix.

technology to be constant, undermining the precise aim of this study. The results are therefore relegated to Appendix 1.B.1.

1.4.2 Results and Counterfactuals

1.4.2.1 Predicting Sectoral Reallocation: the Role of the I-O Change

The results in this section answer the question of how much of the total labor reallocation can be explained by the change in the I-O structure of the economy alone. Figure 1.5a shows the results of the exercise. The variation in the sectoral linkages of the U.S. economy is indeed capable of capturing a sizable amount of the overall labor reallocation across sectors. By omitting all other possible channels, the present accounting model clearly falls short of the actual data, but the predictive power is substantial, considering the simplicity of the exercise. As shown in Table 1.3, the increase in the services share is equal to 10.35 percentage points of total employment until 2002, almost half of the actual change. The result for agriculture is also noteworthy; the sole variation in the I-O linkages accounts for 82% of the total drop in the employment share of this sector.

Table 1.3: Predicted vs. Actual Changes in Employment Shares

Sector	Data	Prediction	Ratio
Agriculture	-3.99	-3.28	82%
Manufacturing	-18.46	-7.07	38%
Services	22.45	10.35	46%

Note: The actual and predicted changes in the employment shares are expressed as percentage points of total employment. The predicted changes are obtained using the proposed Gross Output model. Period: 1948-2002.

Looking at the evolution of the prediction over time, it is evident that it does not increase linearly over time. Even though the changes in I-O linkages drive the result in the right direction, there are other forces that counterbalance this effect. One of these forces is the change in β_j, the sectoral share of value added in gross output; a fall of this share implies that an industry depends more on intermediate inputs from other sectors, hence its overall weight in GDP and in total employment is reduced. For instance, the service sector has experienced a decrease of β_s from 67% to 63%; in particular, this share rose until 1972 to 72% and then fell sharply until 1987. This fact explains why accounting for intermediates does not capture much of the change during the 1972-1987 period. In recent years, the predictive power of the gross output model clearly improves. There are two main reasons for this; first the fall in β_s has been less pronounced, and second it is precisely during this period that the forces that are the focus of this study really take off. In particular outsourcing has increased much more sharply during the second half of the analyzed period, as already shown in Figure 1.4. The share of PBS in total employment rose from 2.2% in 1948 to 11.2% in 2007, but the growth was uneven: 2.8 percentage points accrued between 1948 and 1977, while the increase in the 1977-2007 period was 6.2 percentage points, more than twice as large as the first half.

1.4.2.2 The Rise in Services: the Role of Outsourcing

The other main goal of this study is to quantify the impact of outsourcing on labor reallocation, and on the rise in services in particular. This goal is achieved through the three counterfactual exercises described in Section 1.3.2. Table 1.4 summarizes the results. The overall estimates for the baseline case are again displayed: the current accounting model can explain an increase of 10.35 percentage points in the employment share of services. When the first counterfactual experiment is performed, namely when all I-O coefficients for manufacturing are kept constant to their 1947 level, the prediction drops to 4.01 percentage points, 39% of the value for the baseline case. This result implies that outsourcing could explain 61% of the total prediction obtained in the current framework, in the admittedly far-stretched case that the entire observed change in the shares of intermediate use was coming from outsourcing. Still, this constitutes a useful upper bound for the quantity of interest.

Table 1.4: Effect of Outsourcing on the Service Employment Share

Counterfactual	Predicted Change	Ratio to Baseline	Diff. wrt Baseline
Baseline	10.35	100%	0.00
1: No Service Outsourcing	4.01	39%	6.34
2: No PBS Outsourcing	6.38	62%	3.97
3: No Finance Outsourcing	10.27	99%	0.08

Note: The predicted change and the difference with respect to the baseline setting are expressed in percentage points of total employment. Period: 1948-2002.

Instead when only the PBS share is fixed to its 1947 level, the prediction drops to 6.38 percentage points, 62% of the value for the baseline case. Hence PBS outsourcing accounts for 38% of the prediction generated by the model, which corresponds to an absolute change of 3.97 percentage points, or 18% of the total increase in service employment. This is not a negligible contribution considering that it is considerably more than half of the upper bound and that other types of services are subject to outsourcing, not only PBS. On the other hand, Finance does not seem to contribute much to structural transformation. When the intermediate share of financial services is fixed at its 1947 level, the prediction almost does not move: it drops to 10.27 percentage points, a mere 1% less than in the baseline case.[27]

1.4.2.3 Correcting for the Classification Change

A potential problem with the results presented in the previous section comes from the changes in the classification over time. In fact, while the data for employment and GDP are based on the NAICS classification over the whole period, the data for I-O tables are not. In particular the

[27] In results not shown, I perform another exercise in order to investigate the importance of imported services. The results confirm those already shown by Yuskavage et al. (2006); although the importance of imported services has risen in recent years, their magnitude is still very low, accounting for just 2.7% of total PBS in 2004. Therefore adding non-comparable imports, where most of PBS are concentrated, does not affect the contribution of outsourcing by much.

classification changes in 1997 and, in all previous years, I-O tables are constructed according to the SIC classification. Given that the study is performed at a quite aggregate level, considering three sectors only, most of the changes are not a source of concern because they take place within each sector. Unfortunately there are two major changes that can affect the results: the treatment of publishing and the treatment of auxiliary units. Both were classified within manufacturing under SIC, but they are now classified within services under NAICS; this change causes a jump of the data in 1997. In the case of publishing one might argue that the intrinsic characteristics of the activities in the industry have truly shifted over time, from a pure manufacturing task to a more complex, diversified and service oriented business. Hence, if that was true it would be even more correct not to adjust the data in order to pick up this transformation. In fact, the analysis focuses on the change over the entire period, so it is not really important to determine exactly when this shift took place, and even a gradual change would not invalidate the results. Instead the treatment of auxiliary units is more problematic because, as already noted in Section 1.2.2, they are now classified within PBS, while they were in manufacturing under SIC. Hence the change in the classification of this sector in 1997 may cause problems for the quantification of the contribution of PBS outsourcing.

In order to avoid these issues, I rectify the I-O data after 1997 to keep these two sectors within manufacturing. This adjustment also solves most of the concerns with the measure of outsourcing. Auxiliary units are in fact those establishments dedicated to services within a firm; by excluding this sector, the vast majority of the internal transactions is eliminated. Unfortunately I cannot perform this adjustment in an ideal way. Auxiliary units are classified within sector 55 of NAICS, namely “Management of Companies and Enterprises”. This sector is composed of three sub-sectors: “Corporate, Subsidiary, and Regional Managing Offices” (551114); “Offices of Bank Holding Companies” (551111); and “Offices of Other Holding Companies” (551112). The first sub-sector was moved from manufacturing to PBS but the last two were not, in fact they were already classified within services under SIC as well. The trouble is that I-O data are not disaggregated enough to distinguish these three sub-sectors, hence by re-classifying the entire sector within manufacturing I underpredict the contribution of PBS. In the case of publishing instead, the re-classification can be performed quite precisely, at least for the benchmark years. Finally, the definition of the PBS industry under the two classifications does not match exactly and I have to perform a further finer adjustment within PBS.[28]

Figure 1.5b compares the predictions of the models against the data after the re-classification.[29] As expected, the change in the services share that can be accounted for is lower, but the picture is not so different from before. The change in the I-O structure of the economy is still capable of capturing a sizable amount of the overall labor reallocation across sectors. Also note how the predicted increase in the service share gets smoother over time, reflecting the elimination

[28] See Appendix 1.A.1.1 for the details.

[29] Also the actual data have been adjusted in order to reflect the re-classification of Publishing and auxiliary units. Instead the adjustment within PBS cannot be performed because the industry data are not detailed enough. This introduces a lower bias when the predictions are compared with the data. In fact, when I only exclude the auxiliary units but do not perform the PBS adjustment, the total PBS employment under NAICS is larger than under SIC, as shown in Figure 1.4b. Hence the predictions are compared against employment data that are larger than they should be.

of the problems caused by the change of the classification in 1997. Tables 1.5 and 1.6 report the results. The predicted change in the service share is equal to 7.42 percentage points of total employment, which corresponds to 33% of the actual change. Given that all other channels have been shut down, the prediction is sizable, and it might be a lower bound. The estimate of the drop in the agriculture sector even improves; the changes in the I-O linkages alone account for 86% of the actual drop in agriculture.

Table 1.5: Predicted versus Actual Changes in Employment Shares - No Auxiliaries

Sector	Data	Prediction	Ratio
Agriculture	-3.99	-3.43	86%
Manufacturing	-18.28	-3.99	22%
Services	22.28	7.42	33%

Note: See notes in Table 1.3.

Table 1.6: Effect of Outsourcing on the Service Employment Share - No Auxiliaries

Counterfactual	Predicted Change	Ratio to Baseline	Diff. wrt Baseline
Baseline	7.42	100%	0.00
1: No Service Outsourcing	2.81	38%	4.61
2: No PBS Outsourcing	4.40	59%	3.02
3: No Finance Outsourcing	7.21	97%	0.21

Note: See notes in Table 1.4.

The results on the contribution of outsourcing are also robust. Service outsourcing potentially accounts for 62% of the total prediction; and if the contribution is more plausibly narrowed to PBS only, outsourcing explains 41% of the total, corresponding to 3.02 percentage points of total employment. Given the actual change of 22.3 percentage points, PBS outsourcing alone can explain 14% of the total increase in the share of services in total employment. This share could be subject to a downward bias given the problems with the re-classification and the impossibility of fully adjusting the actual data. Moreover, if the analysis is restricted to 1987-2002, the period in which outsourcing was more pronounced, PBS outsourcing can explain 21% of the total increase in services.

1.5 Final Demand Channels

In this section I allow for for the final uses expenditure shares to vary over time. This exercise proves that the contribution of outsourcing is not wiped out by the traditional final demand channels, and quantifies the extra prediction obtained by accounting for intermediates with respect to a traditional value added model. In the previous section final demand channels were completely shut down, hence a value added model would have simply predicted no reallocation: what the framework in gross output accounted for was essentially an extra prediction with respect to a value added model. In this section, a value added model is capable of predicting a positive labor reallocation through the change in final uses shares, hence the comparison becomes more meaningful. In what follows, I first modify the accounting model to allow for the traditional final demand channels, then calibrate it in this more complicated setting and finally replicate the results of the previous section.

1.5.1 Back to the Accounting Model: Preferences

As pointed out in the introduction, two main channels have been proposed in the structural transformation literature to model the evolution of consumption shares: income effects due to nonhomothetic preferences as in the "utility-based" explanation or substitution effects due to differential productivity growth across sectors as in the "technological" explanation. The main purpose of modeling the evolution of consumption shares is to show that the contribution of sectoral linkages, and of outsourcing in particular, is not negligible even when the standard channels in the literature are present. There is no strong reason to choose one explanation versus the other, but the "technological" approach is adopted here because it is closer to the spirit of this paper and, as pointed out by Ngai and Pissarides (2007), it maintains the independence between parameters of preferences and technologies. Moreover it is more conservative in the number of parameters that need to be estimated, in fact, only the elasticity of substitution is needed while everything else is directly observable.

Consumers take the sector prices P_j as given and maximize their period utility subject to their budget constraint as follows:

$$\max_{\{C_j\}_{j=1}^{J}} \left(\sum_{j=0}^{J} \psi_j C_j^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} \qquad \text{s.t.} \qquad \sum_{j=0}^{J} P_j C_j \leq wL \tag{1.14}$$

where $\psi_j > 0$ and $\sum_{j=0}^{J} \psi_j = 1$. $\epsilon > 0$ represents the elasticity of substitution across sectoral goods. The optimal consumption of each sectoral good is given by:

$$C_j = \frac{\psi_j^{\epsilon} P_j^{-\epsilon} wL}{P^{1-\epsilon}} \tag{1.15}$$

where $P = \left(\sum_{j=0}^{J} \psi_j^{\epsilon} P_j^{1-\epsilon} \right)^{\frac{1}{1-\epsilon}}$ is the aggregate price index. It is possible to define the consumption (or final uses) expenditure share of each sector j as follows:

$$X_j = \frac{P_j C_j}{wL} = \psi_j^{\epsilon} \left(\frac{P_j}{P} \right)^{1-\epsilon} \tag{1.16}$$

To simplify the empirical implementation, let x_j denote the ratio of the consumption expenditure on the good j to the consumption expenditure on the manufacturing good. Re-introducing time subscripts, the new variable is defined as follows:

$$x_{jt} = \frac{X_{jt}}{X_{mt}} = \left(\frac{\psi_j}{\psi_m}\right)^{\epsilon} \left(\frac{P_{jt}}{P_{mt}}\right)^{1-\epsilon} \tag{1.17}$$

And its logarithmic growth rate, $\hat{x}_{jt}$, is simply:

$$\hat{x}_{jt} = ln(x_{jt}) - ln(x_{jt-1}) = (1-\epsilon)(\hat{P}_{jt} - \hat{P}_{mt}) \tag{1.18}$$

Given the absence of capital and investment in this economy, the previous expressions hold for any sector j, including manufacturing for which the growth rate is obviously zero and x_{mt} is always equal to one, at any time t. Exactly as in Ngai and Pissarides (2007), if the elasticity of substitution across composite goods is less than one, the consumption expenditure share expands in sectors with relatively high price growth rates. The opposite holds true if the elasticity is larger than one; and there is no change in consumption shares if the elasticity is exactly equal to one. Given that the sectoral price indexes can be obtained from the data, equation (1.18) is all one needs to get the evolution of the consumption expenditure ratios over time. At each point in time, the sectoral consumption expenditure share, defined in (1.16), can be obtained as follows:

$$X_{jt} = \frac{x_{jt}}{\sum_{k=0}^{J} x_{kt}} \tag{1.19}$$

These shares can then be plugged into equation (1.12) to get the labor shares.

1.5.2 Calibration

When final uses shares are allowed to vary over time, the calibration procedure is a bit more involved. I calculate the final uses ratios relative to manufacturing using (1.17) and their evolution over time using equation (1.18). It is evident from the latter equation that the extra information needed are the sectoral price growth rates and the value of the elasticity of substitution. The latter is set to 0.5, as in Buera and Kaboski (2009). Although there is no final consensus in the literature about the value of this key parameter, $\epsilon = 0.5$ seems a sensible choice given that it is in between the unitary elasticity case often used in the "utility-based" structural transformation literature[30] and the Leontief preferences case ($\epsilon = 0$), which is obtained by Herrendorf et al. (2013b) by minimizing the distance between the expenditure shares predicted by their model and the data. This choice is not far from the value of 0.4 found by Duarte and Restuccia (2010) by matching the share of hours in manufacturing over time and the annualized growth rate of aggregate productivity; and it is slightly smaller than the value of 0.76 found by Acemoglu and Guerrieri (2008) in a two-sector model of high versus low capital intensive industries. Notice that keeping final uses shares constant over time is equivalent to setting the elasticity of substitution to 1. With a unitary elasticity, households use a constant share of their income on each good, and there is no change in final uses shares, as equation (1.18) shows. The results in the

[30]This strand of the literature usually uses "Stone-Geary" preferences, as, for instance, in Caselli and Coleman II (2001) and in Kongsamut et al. (2001) .

previous section precisely correspond to this case.

In order to evaluate the empirical contribution of accounting for intermediates, I compare the results obtained in the proposed gross output framework with those of a benchmark value added model; this is easily obtainable in the present accounting framework by setting $\beta_j = 1$, $\forall j \in J$. When the price channel is shut down as in the previous section, the predictions of the benchmark model are rather humdrum, as it simply predicts no labor reallocation. With less than unitary elasticity and differential price growth rates across sectors, the empirical comparison with the value added model becomes more meaningful. The exercise requires some care, though, as the right set of prices needs to be chosen. For the value added model the choice is quite simple since value added price indexes by industry are immediately available. The sectoral prices indexes provided by the BEA are chain-type annual-weighted indexes, which are not additive. I therefore use the standard methodology for chain price indexes in order to aggregate them up at the three sector level.[31] Figure 1.6 displays the calculated price indexes for the three main sectors; as well-known, when valued added prices are used, services are the sector with the highest increase, followed by manufacturing and then agriculture. These patterns produce changes in the employment shares that are consistent with the stylized facts on structural transformation; according to the model, a higher relative growth in the sectoral price index implies an increase in the consumption share of that sector, and in turn an increase in the employment share.

However, setting a less than unitary elasticity poses extra difficulty when the proposed accounting framework is used. The model is expressed in gross output, hence naively using value added price indexes would not be correct. A first fix would be to use the final consumption expenditure prices, as in Herrendorf et al. (2013b). They use the NIPA tables from the BEA to obtain the price indexes for personal consumption expenditures. They define the three main sectors of interest as follows: agriculture is identified with the NIPA category "food and beverages purchased for off-premises consumption"; manufacturing includes "durable goods" and "nondurable goods" apart from food; services include "services" and "government consumption expenditure". Unfortunately, this approach cannot be adopted in the current framework because it does not match the definition of final uses in the I-O data. A more involved procedure is therefore needed for two main reasons. First of all, the identification of agriculture with the food and beverages category is not correct because it also includes processed products, which are actually produced by the manufacturing sector and hence are classified as manufacturing commodities according to I-O data. Suffice to notice that, in 2002, the expenditures on food and beverage are seven times larger than the personal consumption expenditure associated with agriculture in I-O data. Second, I-O data are in producers' prices while NIPA tables are in purchasers' prices, thus transportation, retail and wholesale margins have to be removed.

I therefore use more disaggregated data to match the I-O final uses to the corresponding NIPA categories, and then transform the series to producers' prices; all the details and data sources are described in Appendix 1.A.3.[32] Figure 1.7 displays the obtained price indexes for final uses

[31] See for instance Whelan (2002).

[32] A further extra adjustment in the price indexes is needed in case investment is also considered. Results for this case are obtained in Appendix 1.B.3.

of the three main sectors; they are also compared to the price indexes used by Herrendorf et al. (2013b). As a robustness exercise, results are also obtained with this alternative set of price indexes; it is already clear from this figure that the predictions will improve considerably in this case. In fact, the price index for services displays a higher growth rate, causing a stronger reallocation. It is also interesting to notice that, in both sets of price indexes and conversely to value added data, final uses prices for agriculture grow more than the corresponding prices for manufacturing.

1.5.3 Results with Variation in Final Uses

The results of the previous section are re-obtained here allowing for the final uses expenditure shares to vary over time. By setting a less than unitary value of the elasticity of substitution, the differential in the price growth across sectors induces a reallocation in the consumption shares, as equation (1.18) shows. In predicting the changes in the sectoral employment shares, the proposed gross output model reacts to changes in the I-O structure of the economy as before. On top of that, both models are now driven by the changes of the sectoral price indexes over time. The results therefore depend on the choice of the price indexes, which have to be constructed in the case of the gross output model. Moreover, the results also hinge on the value of the elasticity of substitution, and hence on the form of consumer preferences. Although this exercise blurs the contribution of all these different channels, it constitutes a good robustness check for the main results of the paper and proves that the contributions of the I-O change and of outsourcing are not wiped out by the standard channel proposed in the literature.

Figure 1.8 plots the predictions of the two models over time, where results have been computed after the re-classifications outlined in Section 1.4.2.3. Given that final uses shares are also allowed to vary over time, the predictions clearly improve but still fall short of the actual data. As shown in Table 1.7, the increase in the services share is equal to 12.98 percentage points of total employment until 2002, which corresponds to 58% of the actual change. If the results (not shown) are computed without performing the re-classification, the share goes up to 69%, which corresponds to 15.45 percentage points of GDP. The overall predictive power also depends on the value of the elasticity of substitution. If one is ready to assume Leontief preferences, the predicted increase in the service share goes up to 17.91 percentage points, 80% of the actual change.

Table 1.7: Predicted vs. Actual Changes in Employment Shares - No Auxiliaries

		Gross Output		Value Added	
Sector	Data	Prediction	Ratio	Prediction	Ratio
Agriculture	-3.99	-3.57	89%	-2.82	71%
Manufacturing	-18.28	-9.42	52%	-5.44	30%
Services	22.28	12.98	58%	8.26	37%

Note: The predicted changes are obtained using both the Gross Output model and the Value Added benchmark model. The elasticity of substitution $\epsilon = 0.5$. See also notes in Table 1.3.

Moreover the results are also affected by the choice of the price indexes. In the current framework, price indexes for final uses are obtained by matching the I-O data to the corresponding NIPA categories and accounting for trade and retail margins. If the personal consumption expenditures indexes proposed by Herrendorf et al. (2013b) are used, the predictions improve considerably. For instance, the predicted change in the service sector employment share rises to 14.82 percentage points, which amounts to 67% of the actual change. As already noted in Figure 1.7, services experience a much higher growth in their price index in this alternative case, and the sectoral reallocation is therefore stronger. Although this alternative set of price indexes is not correct in the current framework, it helps in providing a sense of the robustness of the results with respect to the assumptions I had to take to obtain the preferred set of price indexes. In particular, in order to adjust the prices for the retail and wholesale margins, I have to use value added price indexes instead of the correct gross output prices. This forced choice is likely to have caused a lower bias in the price index for services. In fact, in more recent years, when gross output prices for the retail and wholesale sectors are available, value added prices have experienced a much lower growth compared to gross output prices.[33]

In any case, even if the proposed gross output model cannot perfectly match the data, it is capable of capturing more of the sectoral reallocation compared to the benchmark value added model, over the whole time period. Table 1.7 also shows the predictions obtained with the benchmark value added model. The comparison of the results in the two cases points out that, by accounting for intermediates and allowing for the I-O structure of the economy to change over time, the predictive power is improved. In fact, the extra prediction obtained for the services share amounts to 4.72 percentage points, since the standard model can only predict 37% of the actual change. The prediction for the manufacturing employment share is also much closer to the data: the value added model predicts a drop of just 5.44 percentage points while the proposed gross output model can account for 52% of the total fall, equal to 9.42 percentage points. Finally, it is interesting to note that the prediction is considerably improved in the case of agriculture as well, despite the fact that the gross output price index for agriculture rises more than that for manufacturing; this result once again highlights the importance of the change in sectoral linkages.

Given that more channels are now operating at the same time, I compare the contribution of outsourcing against the portion of the prediction that comes from the change in the I-O structure of the economy. The value of interest is therefore the difference in the predictions of the two models (extra prediction). What is predicted by the benchmark value added model is in fact driven by other channels, like consumer preferences and price changes. Similarly to Tables 1.4 and 1.6, the results of the counterfactual exercises are summarized in Table 1.8; the only difference is that the contribution of outsourcing is now compared against the extra prediction. The overall estimates for the baseline case are again displayed: the current accounting model can account for an increase of 12.98 percentage points in the employment share of services, 4.72

[33]For instance, the value added price for the retail sector, which accounts for most of the margins, experienced a total growth of 14% in the 1987-2007 period; whilst the growth for the gross output price was 28% over the same period. For the wholesale sector the difference is even sharper: the total growth of value added price was just 2% versus a growth of 16% for gross output. See Appendix 1.A.3 for the details on the construction of price indexes.

Table 1.8: Effect of Outsourcing on the Service Employment Share - No Auxiliaries

Counterfactual	Predicted Change		Extra Prediction		
	Value Added	Gross Output	Difference	Ratio to Baseline	Diff. wrt Baseline
Baseline	8.26	12.98	4.72	100%	0.00
1: No Service Outsourcing	8.26	9.16	0.90	19%	3.82
2: No PBS Outsourcing	8.26	10.47	2.20	47%	2.52
3: No Finance Outsourcing	8.26	12.81	4.55	96%	0.17

Note: The Extra Prediction is defined as the difference between the employment share change predicted by the proposed Gross Output model and the change predicted by the Value Added benchmark model. The elasticity of substitution $\epsilon = 0.5$. See also notes in Table 1.4.

percentage points more than the benchmark model. When the first counterfactual experiment is performed, namely all I-O coefficients for manufacturing kept constant to their 1947 level, the extra prediction drops to 0.9 percentage points, 19% of the value for the baseline case. This result implies that, when the price channel is also at work, the difference between the two models is almost entirely captured by the variation in the linkages of the manufacturing sector. Therefore a change in outsourcing policies of manufacturing firms could explain up to 81% of the total extra prediction implied by the current model, in the admittedly far-stretched case that the entire observed change in the shares of intermediate use was coming from outsourcing. Instead when only the PBS share is fixed to its 1947 level, the extra prediction drops to 2.2 percentage points, 47% of the value for the baseline setting; this implies that PBS outsourcing accounts for 53% of the entire extra prediction generated by the model. In absolute terms the contribution of outsourcing amounts to 2.52 percentage points of total employment, slightly lower than the value estimated earlier. Still, this is not a negligible contribution considering that it exceeds 11% of the total increase in service employment and that other types of services are subject to outsourcing, not just PBS.

1.6 Mechanisms of Service Outsourcing

In the simple accounting model proposed in this study, I take the changes in the I-O structure of the economy directly from the data, which corresponds to taking the changes in the parameters of the production functions as exogenous. As firms are changing the mix and the sourcing mode of their inputs over time, an immediate question arises: why is this the case? And in particular, why are firms outsourcing more services over time? A full answer to this question is beyond the scope of this paper, but this section offers some suggestive evidence on the matter, analyzing occupational data and discussing some of the potential drivers.

Outsourcing can take several forms and it is interesting to understand whether firms have: a) outsourced the very same tasks formerly produced in-house; b) substituted inputs produced internally with alternative ones purchased from specialized external suppliers; c) purchased more services from the market in response to new needs. In the first case the change is clearly driven by organizational decisions and represents the starkest form of outsourcing, which would involve a simple relabeling of the same tasks; the mix of activities actually does not change and firms simply outsource what they used to perform in-house. In the second case the firms' choice to outsource might interact with other changes that lead firms to upgrade their activities and outsource them at the same time. In the last case the overall firms' demand for services increases and firms satisfy it through market transactions, rather than internal production.

The results of the previous section apply irrespective of the particular form of outsourcing. Section 1.2.2 showed that the increase in the use of PBS comes from market transactions, and is not matched by a parallel increase in internal production of services. Given the substitutable nature of business services, firms always have the option to employ specialists in-house. If they did not do so there must have been organizational decisions at play. The only potential problem lies with the possibility that the increase in services might be indirectly driven by a change in consumers' tastes. In this particular case organizational changes could be a by-product of a shift in final demand. This section shows that, to a first approximation, the overall composition of activities has not changed over time, and even where specific activities have increased, final demand is unlikely to play a role in that change.

1.6.1 Outsourcing as Relabeling? Evidence from Occupations

Investigating whether firms have outsourced the same tasks they used to produce in-house is an intrinsically difficult exercise because firms' internal activities are very hard to observe (even using data at the firm level). Nevertheless, aggregate occupational data provide some evidence in this regard. In fact, if firms needed more services over time, the occupations involved in the production of these services should become progressively more important, and one should observe an increase of their share in total employment. The challenge is to identify the occupations that best represent the PBS industry. For any given occupation, workers are employed in several sectors and the choice is the result of a trade-off: if only a few occupations are included they will not be representative of the entire PBS industry, but if too many are included the share of workers becomes too large compared to the share of PBS in total employment. I define *PBS Occupations* on the basis of how many workers within each occupation are employed in the PBS industry in 1990. In the baseline definition (Definition 1) I select the occupations that have at least 9% of their workers employed in PBS. As a robustness check, I propose four alternative definitions. Definition 2 uses a threshold of 10%. On the other hand, Definition 3 and Definition 4 are based on the analysis of the PBS industry itself; an occupation is included if at least 0.2% or 0.4% of total workers employed in PBS are classified within that particular occupation. Finally in the "Manual" Definition, I hand pick each occupation on the basis of its job description and whether it could fit in the PBS industry.[34]

Figure 1.9 shows the results of this exercise. Each line plots the share of the selected occupations in total employment over time, according to the different definitions. Interestingly, these shares are fairly constant over time. According to Definition 1, the share of workers classified within the PBS Occupations goes from 24.2% of total employment in 1950 to 28.2% in 2010 but

[34] Data are described in Appendix 1.A.1.2. To obtain consistent occupations over time, the OCC1990 occupational classification scheme is used; occupations are therefore selected using data in 1990.

stays essentially flat from 1970 onwards.[35] It is in this second half of the analyzed period that outsourcing has played a much more important role, as shown in Figure 1.4. In fact, the growth of the share of the PBS industry in total employment in the 1977-2007 period was more than twice as high as in the 1947-1977 period. Therefore PBS increased more sharply in a period when the share of workers classified within PBS Occupations remained constant.

This fact seems to support the idea that what has changed are really the boundaries of the firms, and not so much the underlying activities. Given the rise in the share of the PBS industry in total employment, we expect workers to move from other industries to PBS, or at least the PBS industry to disproportionately employ more workers over time. This is precisely what happens. The share of workers within the selected PBS Occupations that is employed in manufacturing falls over time, while the share that is employed in the PBS industry rises. Figure 1.10a shows the latter share for six main categories used to subdivide PBS Occupations: Managers (and management related occupations); Professionals; Computer related occupations; Clerks, which include various administrative support occupations and some "Service occupations";[36] Technicians; and Other occupations, mainly operators and laborers. Within each category it is evident that the share of workers employed in PBS increases, especially since 1970 when outsourcing really starts taking off. The pattern is particularly sharp for Professionals: the share of workers employed in the PBS industry was 17.5% in 1950, declined to 16.1% in 1970 and has constantly increased since then, reaching 33.2% in 2010 But the growth was even stronger for Technicians and especially for Managers.

Figure 1.11 displays the share of workers employed in PBS for specific occupations. The pattern is quite similar across the board, with a constant increase in this share over time. It is interesting to note that this is true for both high and low skilled occupations. For instance, a very similar growth is experienced by Civil Engineers displayed in panel 1.11c and Guards in panel 1.11d. This fact shows that the rise of PBS is not driven by a particular type of skill and is consistent with both an explanation that focuses on the importance of low-skilled jobs, like in Autor and Dorn (2012), and an explanation that hinges on the rise of high-skilled jobs, like in Buera and Kaboski (2012). At the same time, there are some interesting counter examples. For instance, the share for Lawyers (panel 1.11a) did not change much over time, and was already over 75% in 1950.

The graphical intuition can be more formally established with a standard growth decomposition following Foster et al. (2001). The share of the PBS industry in total employment can be re-written as follows:

$$l_{pbs} = \frac{L_{pbs}}{L} = \sum_{o} \frac{L^{o}_{pbs}}{L^{o}} \frac{L^{o}}{L} = \sum_{o} \omega^{o}_{pbs} l^{o} \tag{1.20}$$

[35] According to Definition 2 and the Manual Definition, the share of PBS Occupations even falls in the second part of the period. The other two definitions are instead a bit more problematic: they include a share of the total work force that is too large. The trade-off between representativeness and over-inclusion becomes clear; Definition 3 includes almost 90% of workers employed in PBS, but at the same time it captures 50% of the total labor force. In the case of Definition 1 the trade-off looks better, in fact it accounts for 82% of the workers employed in PBS but captures just 29% of the total workers.

[36] Note that "Service occupations" is a specific category of the Census Bureau classification and should not be confused with the service sector; it mainly includes low-skilled jobs like Guards, Janitors, and Cleaners, but also mid-skilled jobs like Dental assistants and Health aides. See Appendix 1.A.1.2 for precise definitions.

where ω^o_{pbs} represents for a given occupation o the share of workers that are employed in the PBS industry (displayed in Figure 1.10a and Figure 1.11), and l^o is the share of occupation o in total employment. The change in the PBS employment share becomes:

$$\Delta l_{pbs} = \underbrace{\sum_o \Delta\omega^o_{pbs} l^o_1}_{Within} + \underbrace{\sum_o \omega^o_{pbs,1} \Delta l^o}_{Between} + \underbrace{\sum_o \Delta\omega^o_{pbs} \Delta l^o}_{Cross} \tag{1.21}$$

where l^o_1 and $\omega^o_{pbs,1}$ indicate quantities at the beginning of the period. The first term is a within-occupation component that captures how much of the increase in PBS employment is due to workers within each occupation moving to the PBS industry, while the second term is a between-occupation component that captures the contribution of employment share reallocations among occupations. I perform the decomposition for the 1970-2010 period and split occupations according to the main categories previously introduced, plus an extra category that includes all other occupations not classified as PBS Occupations.

Table 1.9: Decomposition of the PBS Employment Share Growth

Category	Within	Between	Cross	Total
Managers	0.86	0.01	0.05	0.92
Professionals	1.04	0.16	0.17	1.36
Computer	0.11	0.27	0.38	0.76
Clerks	1.39	-0.16	-0.18	1.05
Technicians	0.17	-0.04	-0.03	0.10
Others	0.05	-0.05	-0.02	-0.02
Not-PBS Occupations	1.52	-0.01	-0.01	1.50
Total	5.14	0.19	0.34	**5.67**

Note: The grand total (in bold) is the increase in the PBS industry share in total employment over the 1970-2010 period, all numbers are in percentage points of total employment. Data from IPUMS-USA, unemployed and workers with unknown occupation or industry are excluded.

Table 1.9 reports the results of the decomposition. Most of the growth comes from the within component: workers do not change occupation but move to PBS from other industries, mainly manufacturing (or are disproportionately more likely to be hired in PBS). The between component accounts for a very marginal share of the total growth, so the rise or fall of certain types of occupations does not account for much of the increase in PBS employment, which supports the idea that the underlying activities have remained roughly constant over time. The same result holds true for almost all categories. The main exception is Computer related occupations for which the between and the cross components play a bigger role, but this is intuitive given that this type of occupation did not exist before 1970. Although marginal in relative terms, the between and cross components for Professionals and Clerks are smaller but comparable to those for Computers in absolute terms. So, to further investigate the role of the reallocation of employment shares among occupations, Figure 1.10b shows the breakdown of PBS Occupations into their main categories, where, for each category, I plot its share in total employment (l^o). Despite the total share of PBS Occupations being roughly constant over time, there is some heterogeneity across categories, as partially revealed by the decomposition. In particular, the share of Clerks falls when the share of Computer related occupations rises.

This pattern provides suggestive evidence about other changes that occurred over the period. Outsourcing might not take place through the mere substitution of the very same task from inside to outside the firm, but it could entail the substitution of an old superseded task with a new, more technologically advanced one. In this sense, outsourcing could be a way of accessing new technologies that would be too costly to produce in-house, as Bartel et al. (2012) have argued. The substitution of computer specialists employed in specialized service firms for clerks employed internally is a fitting example. At the same time, the share of Professionals also rises over time, suggesting an increase in the need of specialized knowledge. The next section discusses these two potential drivers of outsourcing.

1.6.2 Determinants of Service Outsourcing

So why have firms outsourced more services over time? The answer is likely to be related to two intertwined changes. The first comes from the service supply side and consists in the rise of an external market for PBS. Over time more and more firms have specialized in services, and slowly best practices have been established. As argued by Deblaere and Osborne (2010), services have been broken into their components and optimized by eliminating redundancies, automating and standardizing wherever possible. Essentially the production of services has been industrialized, creating a proper market for them, and economies of scale have allowed external providers to beat internal production. This explanation is formalized by Garicano and Rossi-Hansberg (2012a) in a model of growth where organizations develop to exploit existing technologies. They model the process through the emergence of markets for specialized services that are slowly created to satisfy the demand of agents that, facing some exceptional problems, do not have the incentive to acquire the specialized expertise to solve them. The creation of these referral markets takes time because experts have to learn the problems and invest in the knowledge to solve them. The high share of lawyers employed in the PBS industry over the entire period is suggestive in this regard. Law firms have a long history in the U.S. and were already well established in 1950; as a result most lawyers were employed within PBS at the beginning of the period. This shows how the decision of outsourcing services is very much related to the existence of external providers, that is, a market that can provide the services at a given price.

Service outsourcing as a way to access the external provider's specialized skills was first proposed by Abraham and Taylor (1996). The intuition again comes from the fact that it might not be optimal for a firm to invest in these competencies while an external provider can enjoy economies of scale and amortize the sunk costs of these investments across several clients. Although focused on parts and components rather than service outsourcing, Bartel et al. (2012) build on the same intuition to provide a model in which the probability of outsourcing production is positively related to the firm's expectation of technological change. Investing in a new technology implies some sunk costs; the faster technological change, the shorter the lifespan of a new technology, and firms have less time to amortize the sunk costs. Therefore firms outsource to avoid these costs and substitute the old technology with the latest version provided by external suppliers, which can enjoy economies of scale and spread those costs over a larger demand.

On the other hand, from the service intermediate demand side, manufacturing firms constantly strive to grow to increase their scale and profits. The problem is that growing is painful and comes at a cost, for instance, in terms of coordination across business units. Outsourcing has helped firms to grow, allowing them to focus on their core competencies and externalizing the tasks that were not a source of competitive advantage. In essence, outsourcing has been a way to support a more complex environment. In an ongoing research project, I investigate the firm's demand side and build a model of the boundary of the firm based on adaptation costs and diminishing return to management. I look at one possible driver of managerial/coordination complexity: the internationalization decision of the firm. In doing so, I unveil new systematic evidence about domestic service outsourcing. For a large panel of French firms, I find that the share of purchased business services in total costs is positively and significantly related to the number of export destination countries and to the number of products.[37]

A full empirical investigation of the determinants of outsourcing is difficult because firm-level data are not available for the long period of the present analysis. However, interesting insights can be obtained from industry level data over the second part of the period. In particular, I test whether coordination complexity and the need for accessing external skills and new technologies are drivers of service outsourcing. I capture coordination complexity with the complexity of the division of labor, as proposed by Michaels (2007). Specifically complexity of an industry is measured as one minus the Herfindahl index of the occupations of its workers, excluding managers (but results do not change if managers are included). In the absence of exogenous variation in the main variables, it is not possible to give a full causal interpretation of the results. The results are nevertheless informative, and robust to the inclusion of industry fixed effects, year fixed effects and other potential drivers of outsourcing.

I run the following reduced form regression:

$$OUT_{it} = \beta_1 C_{it} + \beta_2 P_{it} + \boldsymbol{W'_{it}\beta_3} + \delta_i + \delta_t + \epsilon_{it} \tag{1.22}$$

where OUT_{it} is the share of purchased business services over total sales for industry i at time t, C_{it} is the complexity of industry i, P_{it} is the number of patents used by industry i, $\boldsymbol{W_{it}}$ is a vector of controls, and δ_i and δ_t are industry and time fixed effects, respectively. I take the measure of outsourcing from I-O tables, where I exclude auxiliary units as in Section 1.4.2.3. Outsourcing is defined as the share of PBS inputs over total sales (direct requirement coefficient) and the industries are defined at the 4-digit SIC level. The analysis is restricted to the manufacturing sector and the data are from the benchmark years 1972, 1982, 1992 and 2002.[38] Occupational data are from the IPUMS-USA database and I use the variable IND1990 to get a consistent

[37] See Chapter 2.

[38] The concordance table created to obtain a consistent definition of SIC industries over time is available on request.

definition of industries over time.[39] Following Bartel et al. (2012), I proxy technological change as the number of patents used by an industry. Patents data according to the International Patent Classification come from the NBER U.S. Patent database (updated version), described in Hall et al. (2001) and available from 1976 onwards. I obtain the number of patents used by an industry (as opposed to patents created by an industry) using the concordance table provided by Silverman (2002).

Table 1.10: Determinants of PBS Outsourcing

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Complexity	2.850^a	5.766^a	5.614^a	5.643^a	6.604^a	6.492^a	6.487^a
	(0.585)	(1.183)	(1.193)	(1.278)	(0.965)	(1.010)	(1.010)
Num. Patents			0.270^b	0.276^b	0.259^b	0.254^c	0.256^c
			(0.128)	(0.129)	(0.130)	(0.134)	(0.133)
Num. Inputs				0.150	0.187^c	0.185^c	0.185^c
				(0.098)	(0.095)	(0.096)	(0.096)
K/L					0.050	0.045	0.051
					(0.064)	(0.065)	(0.069)
S/L						0.045	0.044
						(0.099)	(0.099)
Scale							-0.014
							(0.065)
Observations	1,789	1,340	1,338	1,338	1,329	1,329	1,329
Number of ind.	459	459	458	458	458	458	458
R-squared Within	0.294	0.267	0.276	0.280	0.283	0.283	0.283
Fixed effects	ind&year	ind&year	ind&year	ind&year	ind&year	ind&year	ind&year

Note: The dependent variable is the share of PBS over total sales (direct requirement coefficient). All variables are in logs. Data in column (1) are for years 1972, 1982, 1992 and 2002; in the remaining columns year 1972 is dropped because the number of patents is not available in that year. Industry-clustered standard errors are in parentheses; (a, b, c) indicate 1, 5, and 10 percent significance levels.

Table 1.10 shows the results. Controlling for industry and year fixed effects, coordination complexity is positively and significantly related to service outsourcing. The effect has strengthened over time: from column (2) onwards year 1972 is dropped and the magnitude is higher. The need to access external skills and new technologies, measured as the number of patents used by the industry, also has a positive effect, but it is less robust to the inclusion of year fixed effects and other controls. As an alternative measure of complexity, I also include the number of inputs, or more precisely the share of the number of commodities that the industry uses over the total available commodities (to control for changes in the classification over time). As expected, the impact is positive but only marginally significant after including year fixed effects.

The results are confirmed when other determinants of outsourcing are included. In particular, I add capital intensity, human-capital intensity, and a measure of scale economies at the plant level, as proposed by Antràs (2003).[40] None of the controls have a significant effect in the case of service outsourcing. In Appendix 1.C, I also test the robustness of the findings to an alternative

[39] The concordance table from IND1990 to SIC is available on request. Occupational data are available every ten years, so I measure complexity with a 2-year lead with respect to outsourcing. I do not use data before 1970 because I would lose 25% of the industries.

[40] The data come from the NBER Manufacturing Industry Productivity Database; the number of establishments used to calculate the scale variable is from the County Business Patterns of the U.S. Census Bureau.

measure of service outsourcing taken from the Census of Manufacturing, which avoids all the issues with internal transactions. This alternative data source also allows me to test other determinants of outsourcing as proposed by Yeaple (2006), Nunn (2007), and Costinot et al. (2011), but the data are available from 1992 only. The picture is very similar and both measures of complexity are positively and significantly related to service outsourcing.

The evidence shown in this section supports the view that the overall composition of firms' activities has remained roughly constant over time. And even if few specific activities and occupations have increased their importance over time, the mechanisms at play seem to be related to technology or other supply side channels. Further research at a more micro level is needed but the analysis so far shows that final demand does not play a major role, even an indirect one, in the rise of PBS, an industry that accounts for almost half of the total rise of the service sector.

1.7 Conclusions

By presenting a simple gross output accounting model that can capture the fully-fledged input-output structure of the economy, this paper investigates the role played by firms in shaping the reallocation of resources across sectors. In doing so, it contributes to the structural transformation literature by shifting the focus to forces that drive the process of structural transformation but that, at the same time, are completely unrelated to consumer preferences, namely the choice of the input mix and sourcing mode.

I use the gross output accounting model to evaluate the sectoral reallocation of employment in the U.S. over the period 1948-2002. When both the standard channels in the literature and the forces proposed in this study are at play, the predicted change in the service share is equal to 13 percentage points of total employment in the baseline estimates. This prediction amounts to 58% of the actual change, and is larger than the 37% estimated by a benchmark value added model. When the channels proposed in the literature are shut down by keeping the final uses expenditure shares constant over time, the sole evolution of the input-output structure of the economy can explain a change in the service share equal to 7.4 percentage points of total employment, 33% of the actual change. I perform a counterfactual experiment in order to quantify the contribution of professional and business services outsourcing to the sectoral reallocation. In the same specification, this particular type of outsourcing explains 41% of the prediction, which amounts to 3 percentage points. Given the actual change of 22.3 percentage points, professional and business services outsourcing alone accounts for 14% of the total increase in the share of services in total employment. Interestingly, this estimate is not too far from the back-of-the-envelope calculation performed by Fuchs (1968) over forty years ago. In fact, he showed that the growth of intermediate demand for services by goods-producing industries accounted for less than 10% of the total employment change between 1947 and 1958. The fact that professional and business services outsourcing alone now accounts for more than 10% of the total increase in services can be explained in light of the remarkable increase of this phenomenon in more recent years.

Further research at a more micro level is needed to understand why firms have been outsourcing more services over time. In its starkest form, outsourcing can be interpreted as a mere relabeling of economic activity and the constant share of business services occupations in total employment over time supports this view. But relabeling is not the only interpretation and there is some variation at a more disaggregated level. Under alternative views, outsourcing can be seen as a way to access new technologies or support a more complex business environment, helping firms to focus on their core competencies by externalizing the tasks that are not a source of competitive advantage. A general message of this paper is that more attention should be devoted to services, since future growth will depend more and more on this key sector.

Figures

Figure 1.1: Service Sector Growth in the U.S.

(a) Share of GDP

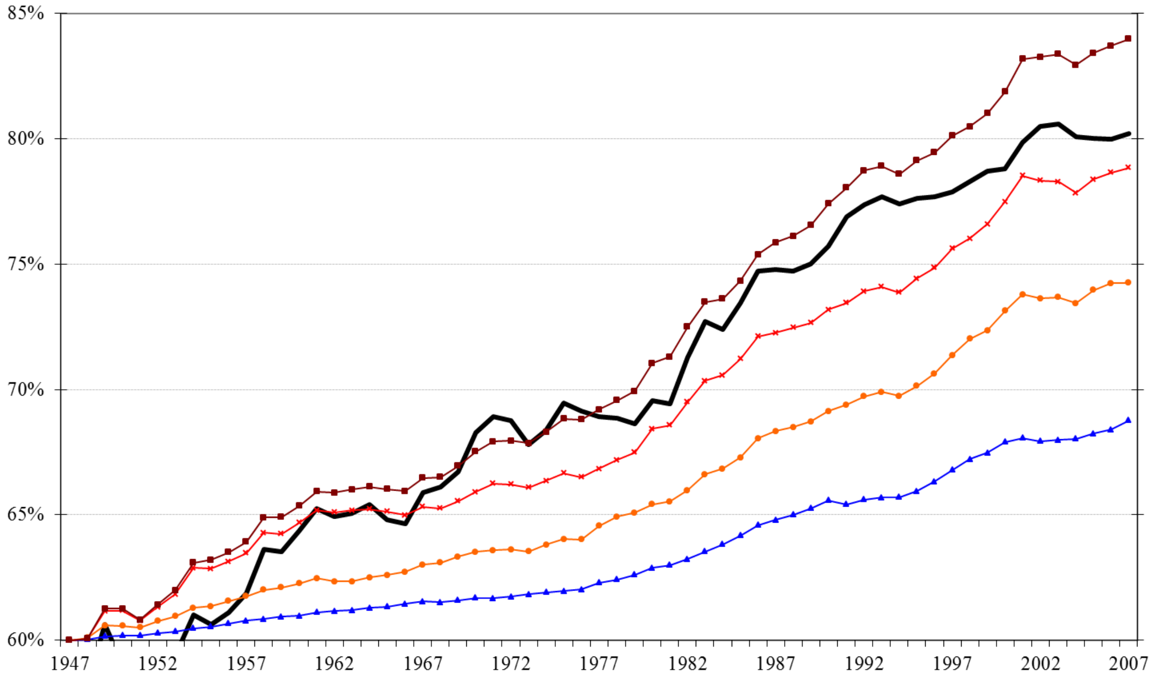

(b) Share of total employment

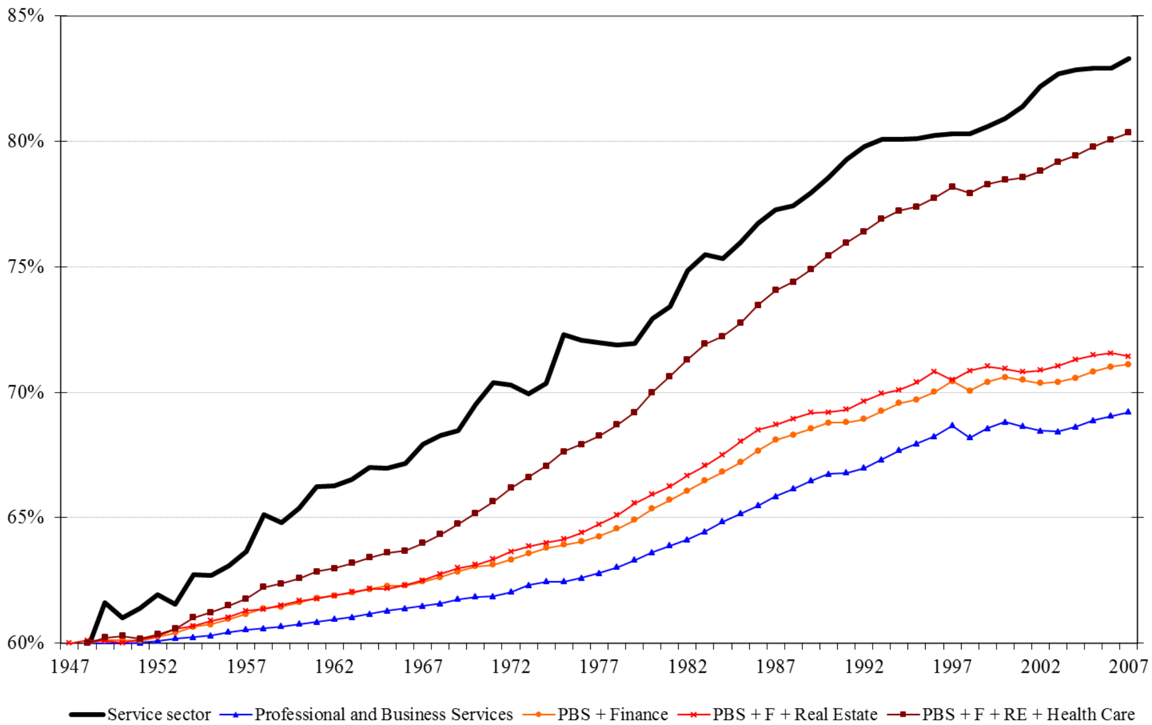

Source: BEA Annual Industry Accounts, release: December 2010.
Note: The left-hand side axis displays the absolute share of the entire service sector (thick black line) in terms of either GPD (panel a) or total employment (panel b). The right-hand side axis applies to all series and displays the change in percentage points of either GPD or total employment. The triangle marked line represents the percentage point change of Professional and Business Services (PBS); the circle marked line represents the percentage point change of the combined sector PBS and Finance; analogously the cross marked line for the combined sector PBS, Finance and Real Estate, and the square marked line for the combined sector PBS, Finance, Real Estate and Health Care.

Figure 1.2: Total Requirements Tables in the U.S., 1947-2002

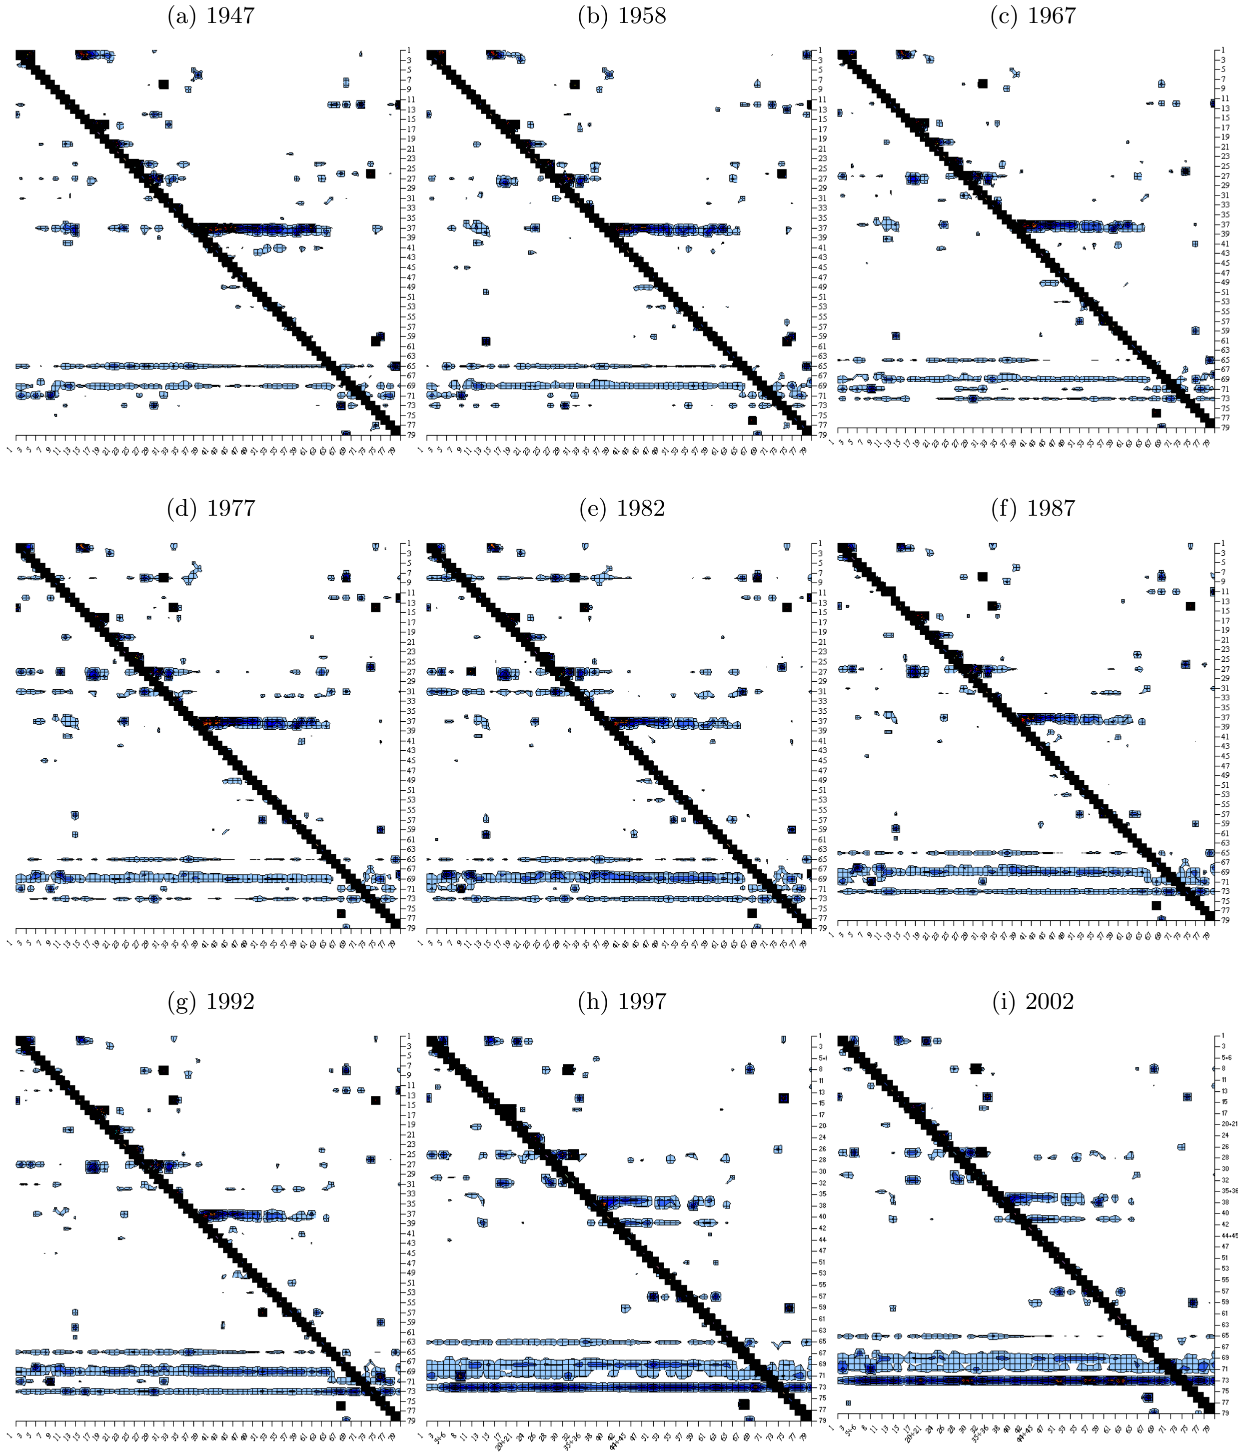

Source: BEA Benchmark Input-Output Accounts.
Note: The tables for years 1947 to 1967 show the 85-industry level total requirements coefficients, the tables for years 1972 to 1982 show the 85-industry level IxC total requirements coefficients; all data are readily available on the BEA website. The tables for years 1987 and 1992 are obtained from the Use and Make tables at the six-digit level. The tables for years 1997 and 2002 are obtained from the Use and Make tables at the summary level and transformed into I-O SIC codes using a concordance table available on request. A contour plot method is used, showing only shares greater than 2% of the total output multiplier (or backward linkage).

Figure 1.3: The Influence of PBS on the U.S. Economy

(a) PBS Total Use

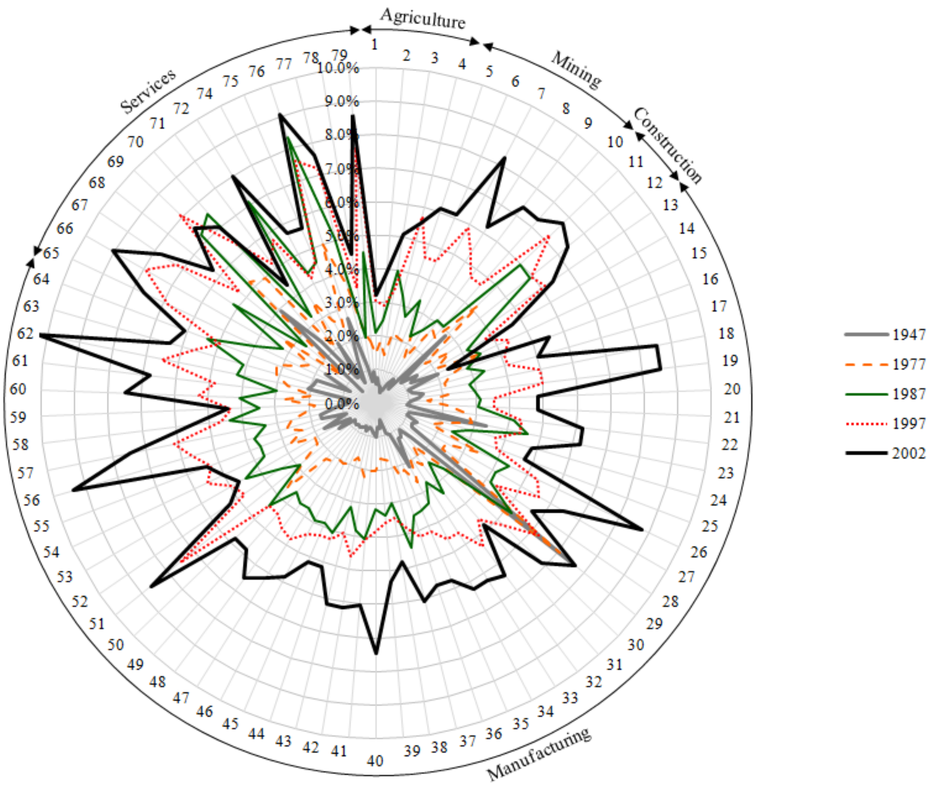

(b) Influence Vector

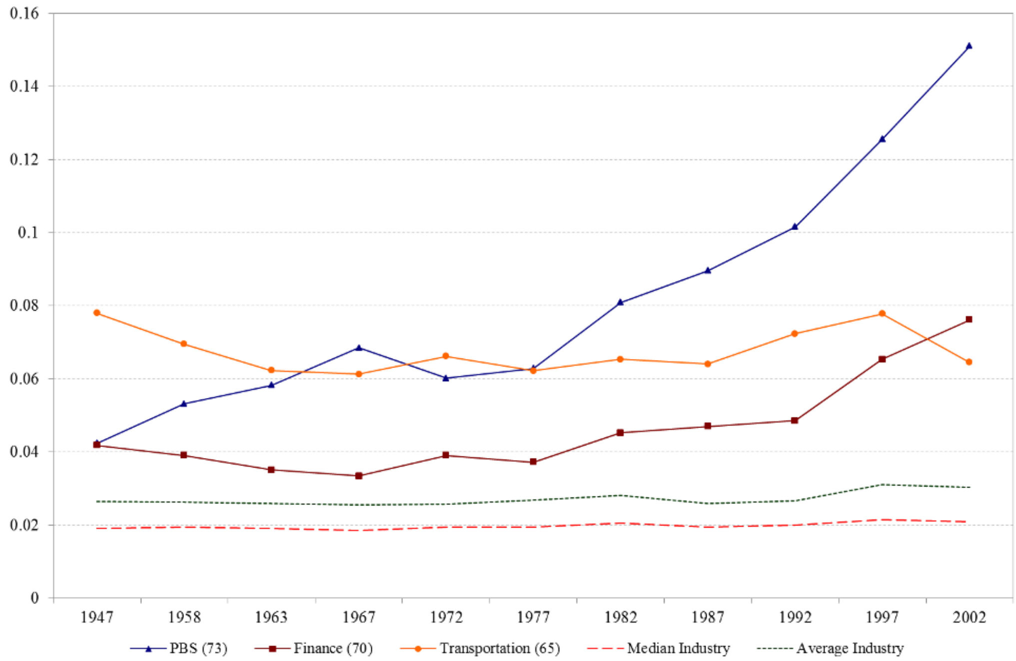

Source: BEA Benchmark Input-Output Accounts and author calculations.
Note: Panel (a) displays the share of PBS in the total requirements for all commodities in the economy (one outlier - Radio and television broadcasting, 67 - is excluded in 1947 for graphical reasons). The influence vector is defined as: $v = \frac{1}{J}\Omega^{-1}\mathbf{1}$, where J is the number of sectors and Ω^{-1} is the total requirements table (see Section 1.3). Panel (b) plots over time the elements of the vector v corresponding to PBS, Finance, Transportation, and the average and the median industry. Auxiliary units are excluded; see Section 1.4.2.3.

Figure 1.4: Professional and Business Services (Share of)

(a) Published Series

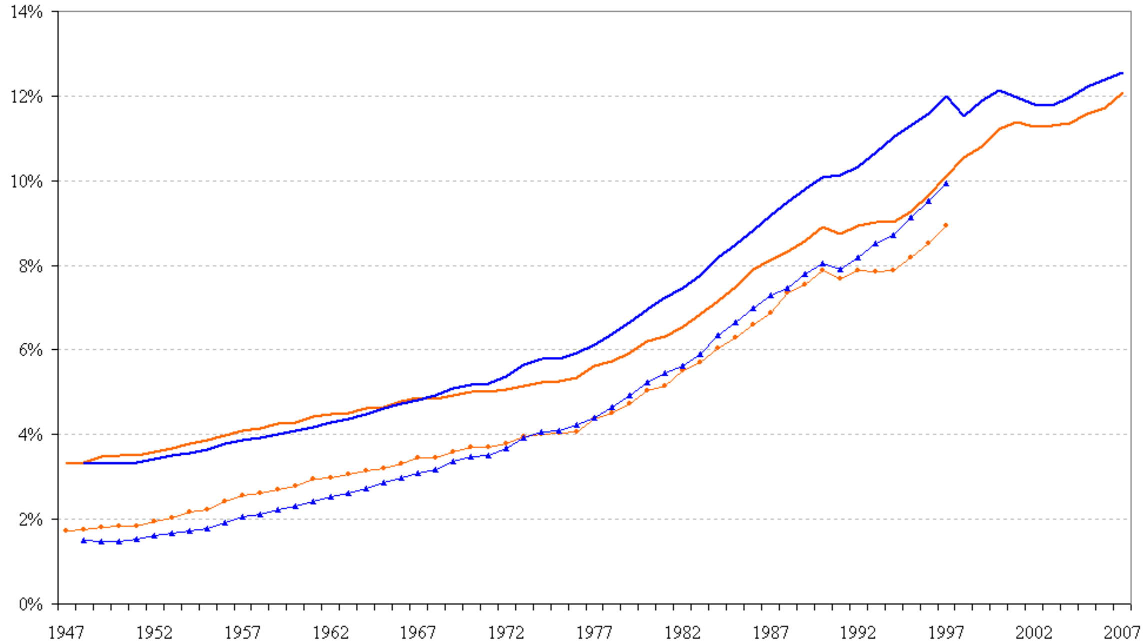

(b) No Auxiliaries

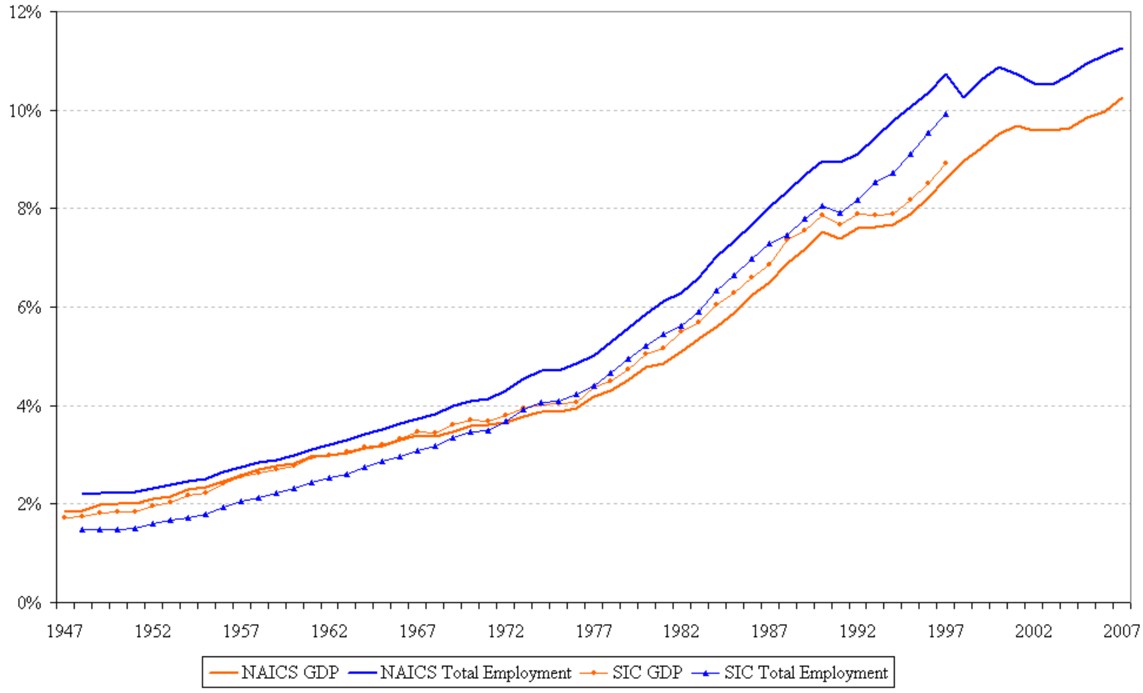

Source: BEA Annual Industry Accounts, release: December 2010.
Note: Professional and Business Services under the 1987 SIC classification include: Business Services (73); Miscellaneous Repair Services (76); Legal Services (81); Other Services (84, 87, 89). The series is not entirely consistent over time; before 1987 the 1972 SIC classification is used, the two coincide apart from Other Services that is named Miscellaneous Professional Services and the corresponding 1972 codes are 84 and 89. Under NAICS Professional and Business Services include: Professional, Scientific, and Technical Services (54); Management of Companies and Enterprises (55); Administrative and Waste Management Services (56). Management of Companies and Enterprises (55) mostly coincide with the so-called auxiliary units under the SIC classification and it has been excluded from the data of panel 1.4b.

Figure 1.5: Predicted vs. Actual Employment Shares in the U.S.

(a) Published I-O Tables

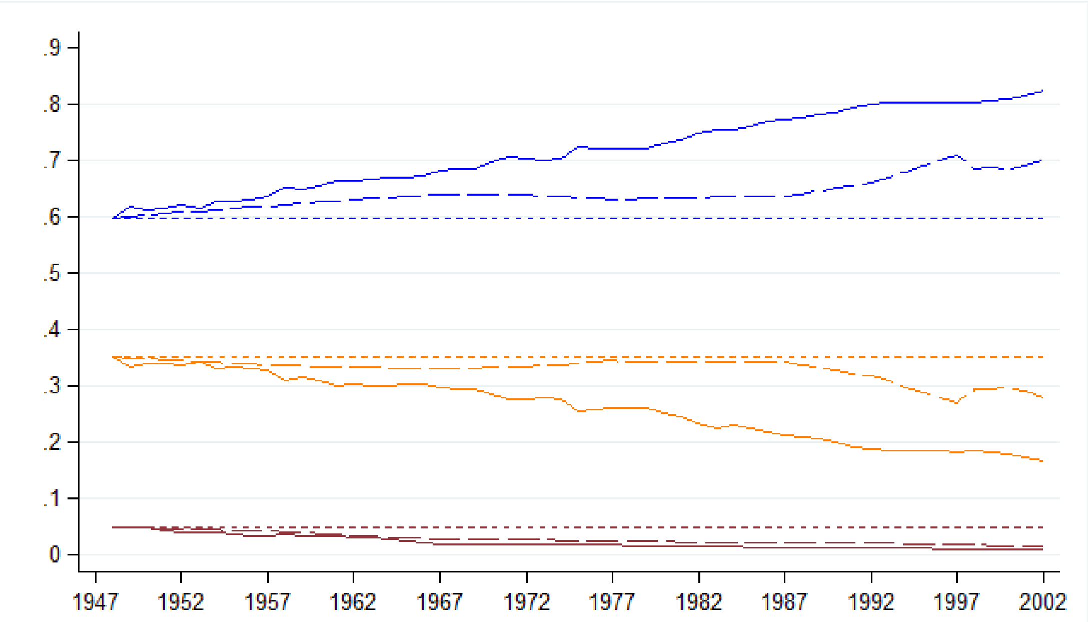

(b) No Auxiliaries

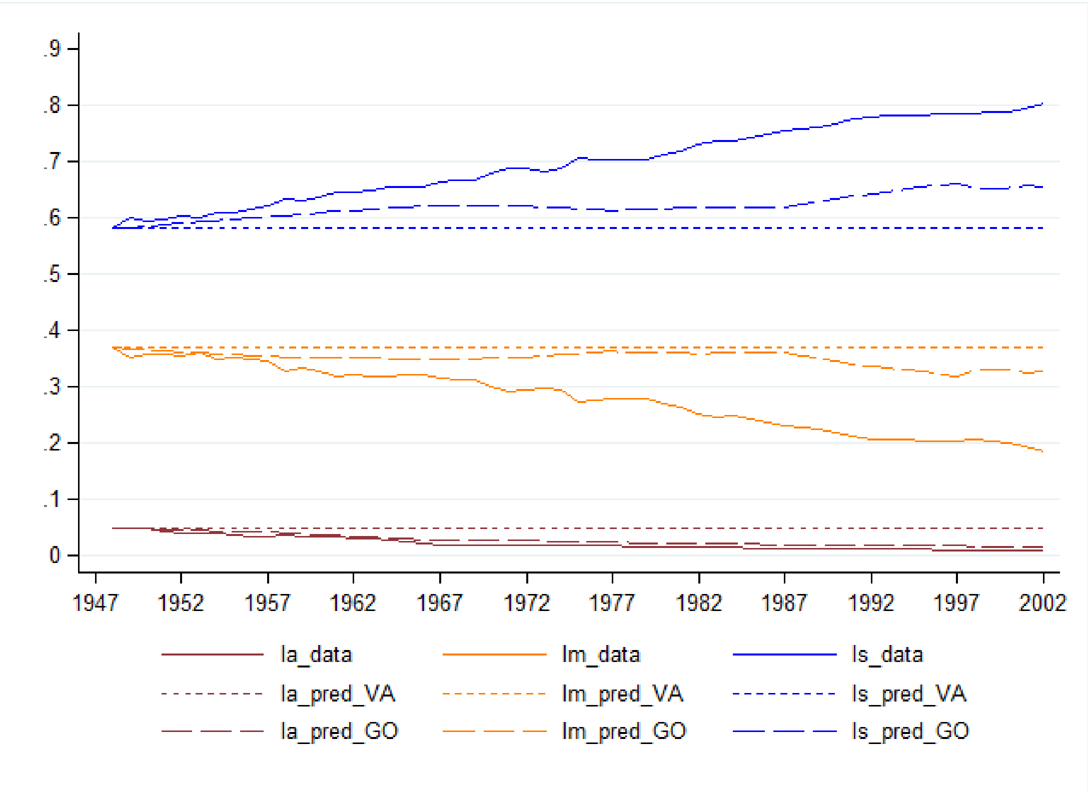

Source: BEA Benchmark and Annual Industry Accounts (release: December 2010) and author's calculations.
Note: Period: 1948-2002. The first panel shows data and predictions obtained using the published I-O tables; the second panel instead is obtained after the re-classification of auxiliary units, PBS and publishing performed in Section 1.4.2.3. The predicted changes in labor shares for agriculture (la), manufacturing (lm) and services (ls) are obtained using the proposed Gross Output (GO) model. A traditional Value Added (VA) model predicts no change because final uses are kept constant.

Figure 1.6: Value Added Price Indexes (1947=1)

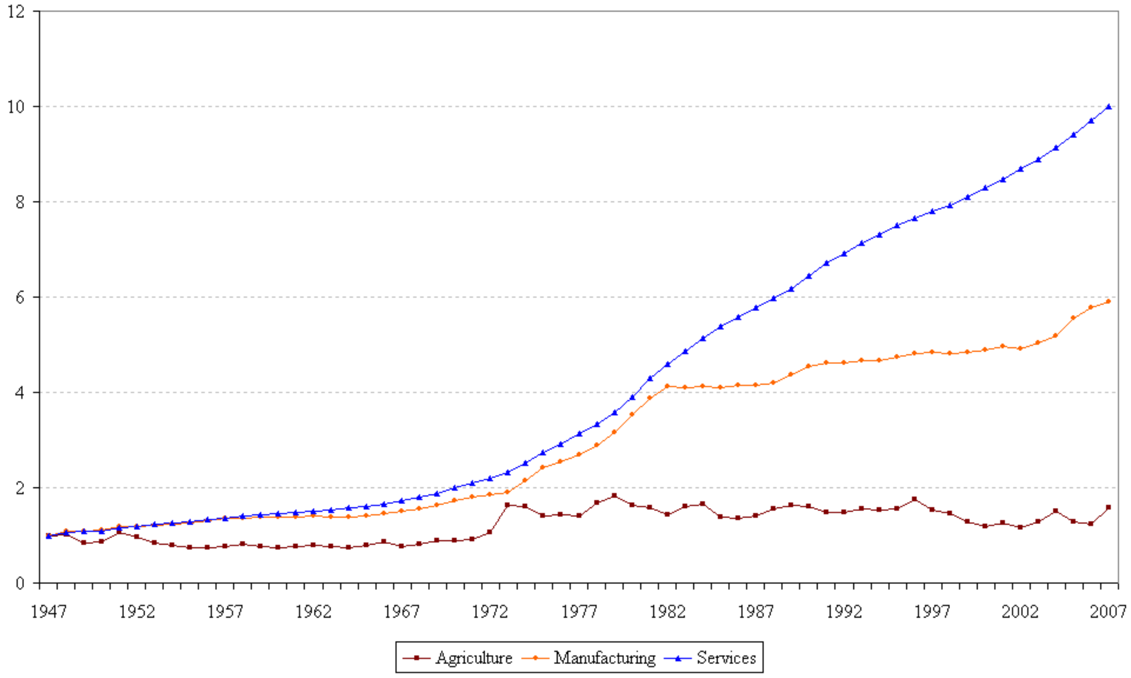

Source: BEA Annual Industry Accounts (release: December 2010) and author's calculations.

Figure 1.7: Final Uses Price Indexes (1947=1)

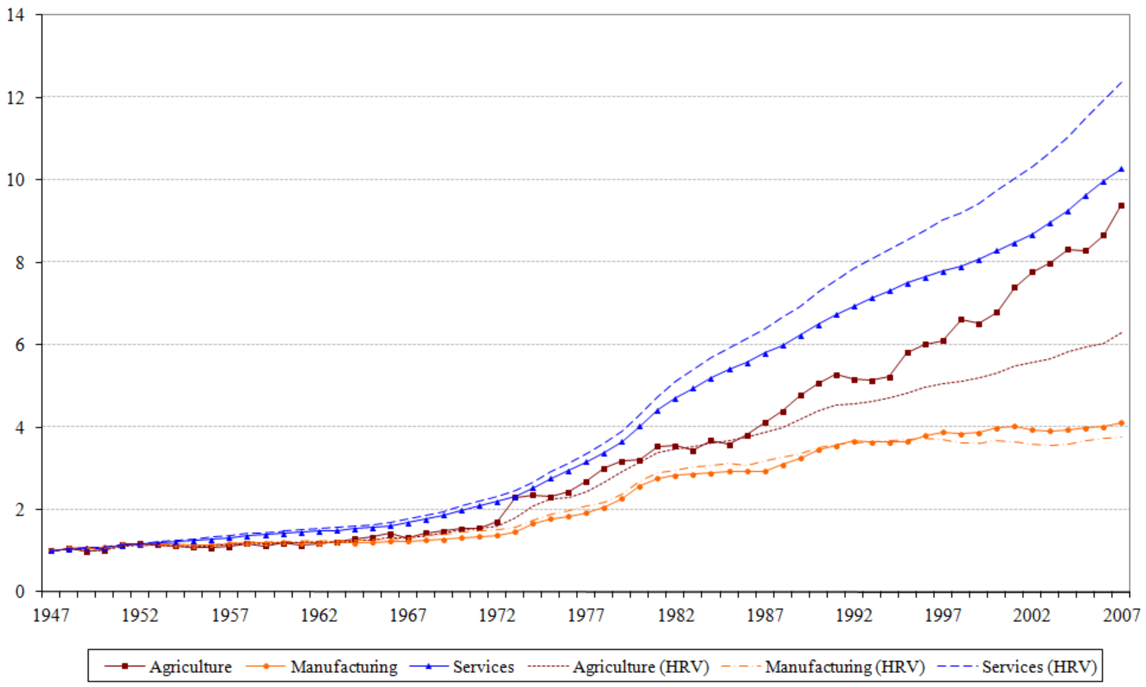

Source: BEA NIPA Tables and author's calculations.

Figure 1.8: Predicted vs. Actual Employment Shares in the U.S.

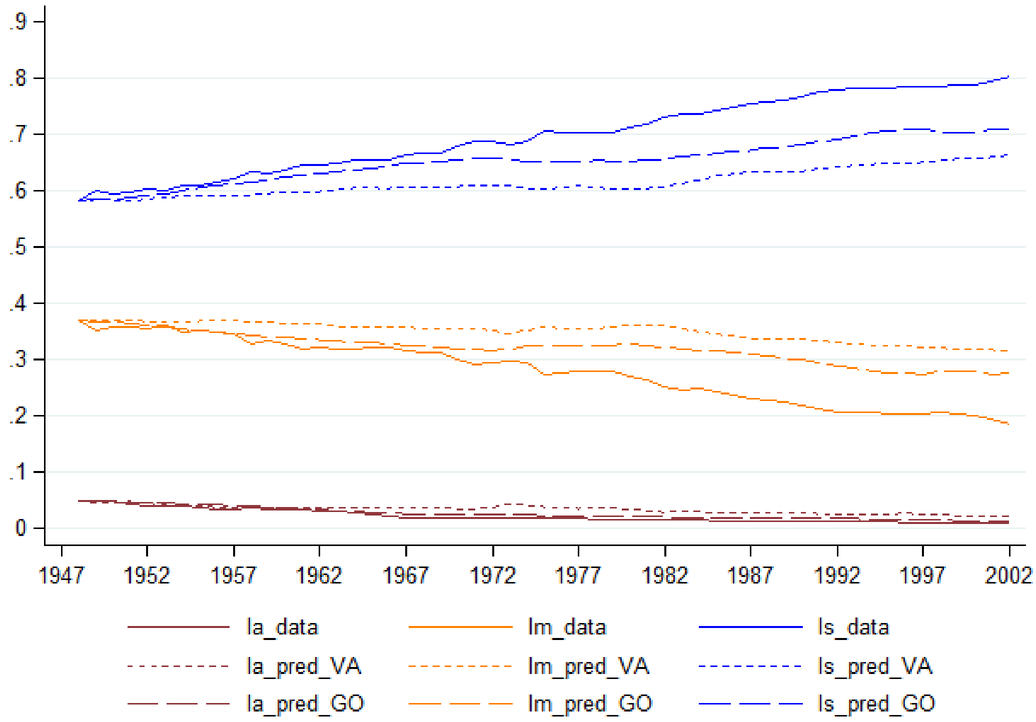

Source: BEA Benchmark and Annual Industry Accounts (release: December 2010) and author's calculations.
Note: Period: 1948-2002. The predicted changes in labor shares for agriculture (la), manufacturing (lm) and services (ls) are obtained using both the proposed Gross Output (GO) model and the Value Added (VA) benchmark model. The elasticity of substitution $\epsilon = 0.5$.

Figure 1.9: Share of PBS Occupations in Total Employment

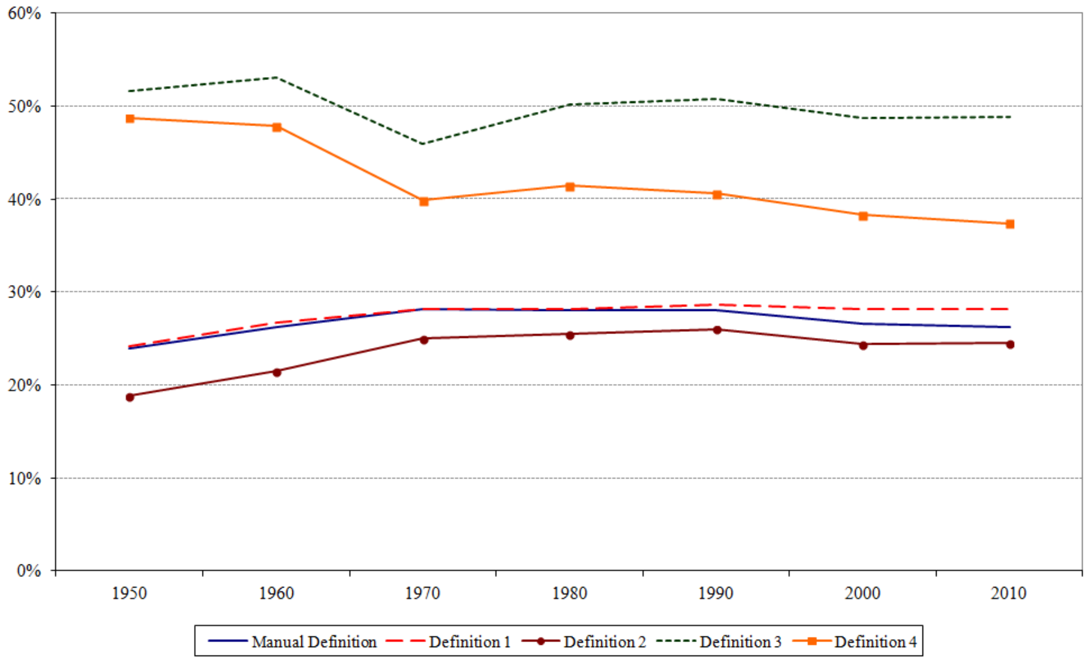

Source: IPUMS-USA.
Note: PBS Occupations are selected according to five definitions, as described in the main text.

Figure 1.10: Main Categories of PBS Occupations

(a) Participation in PBS (Within component)

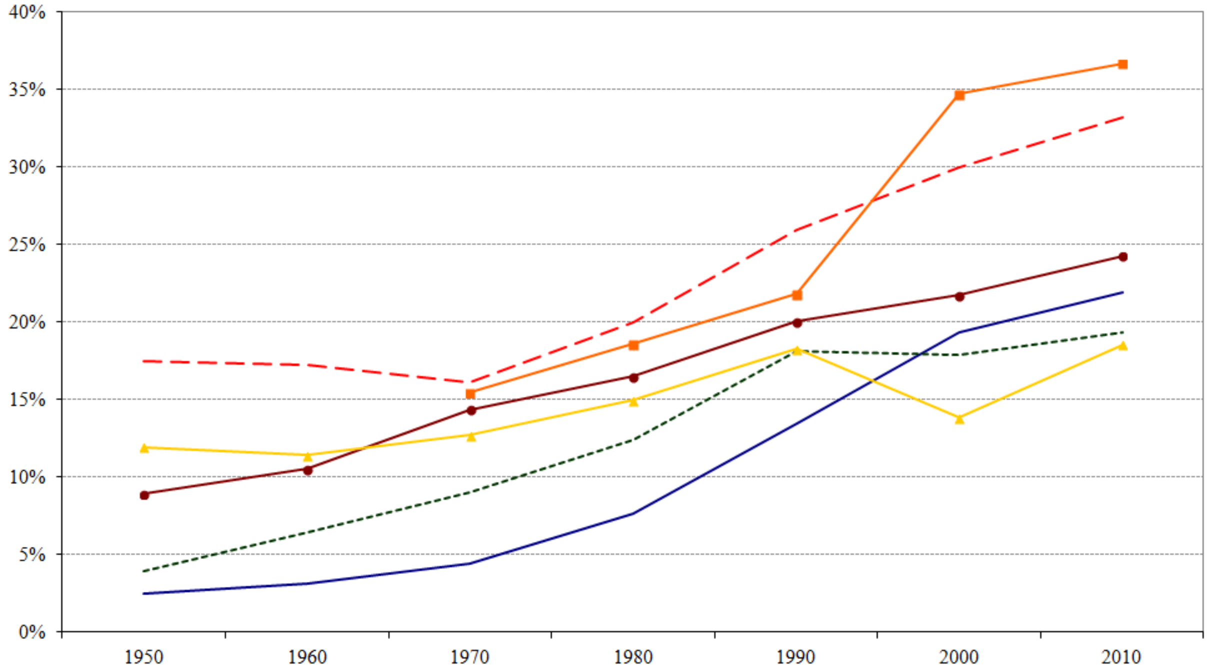

(b) Share in Total Employment (Between component)

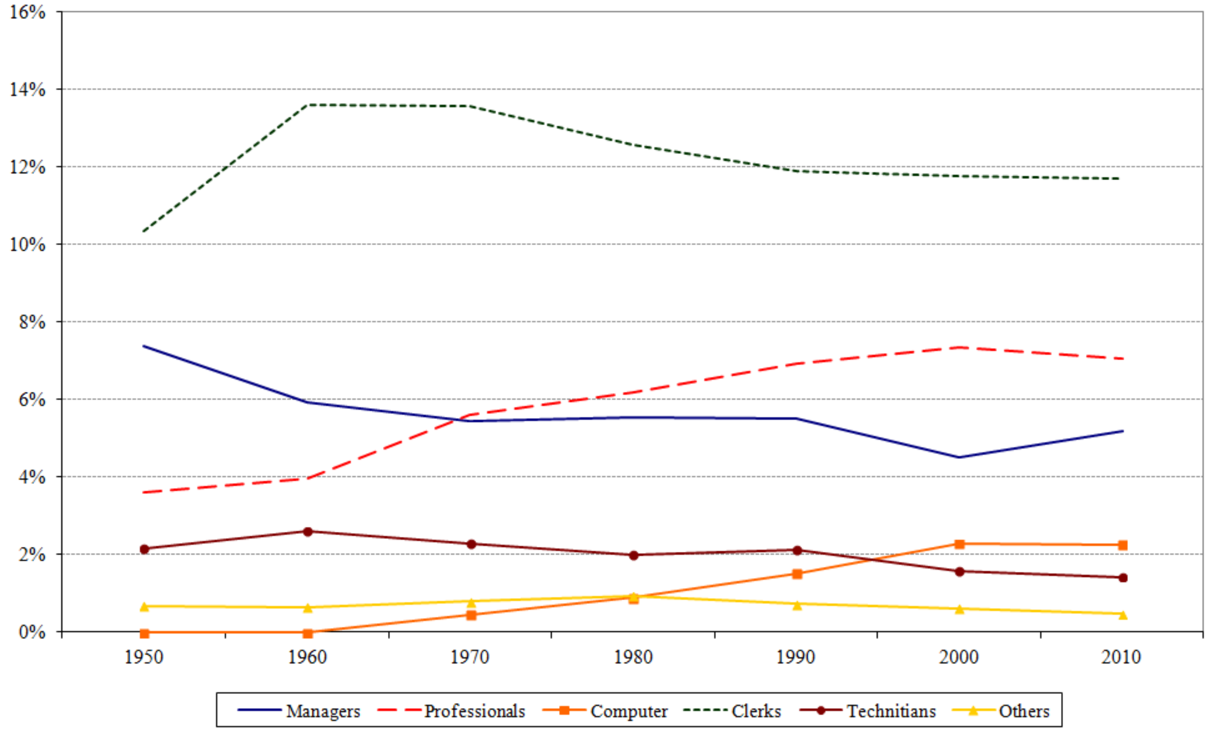

Source: IPUMS-USA.
Note: PBS Occupations are selected according to Definition 1. Panel (a) plots, within each main category, the share of workers that are employed in the PBS industry. Panel (b) plots the share of the main categories in total employment.

Figure 1.11: Selected Occupations - Participation in PBS

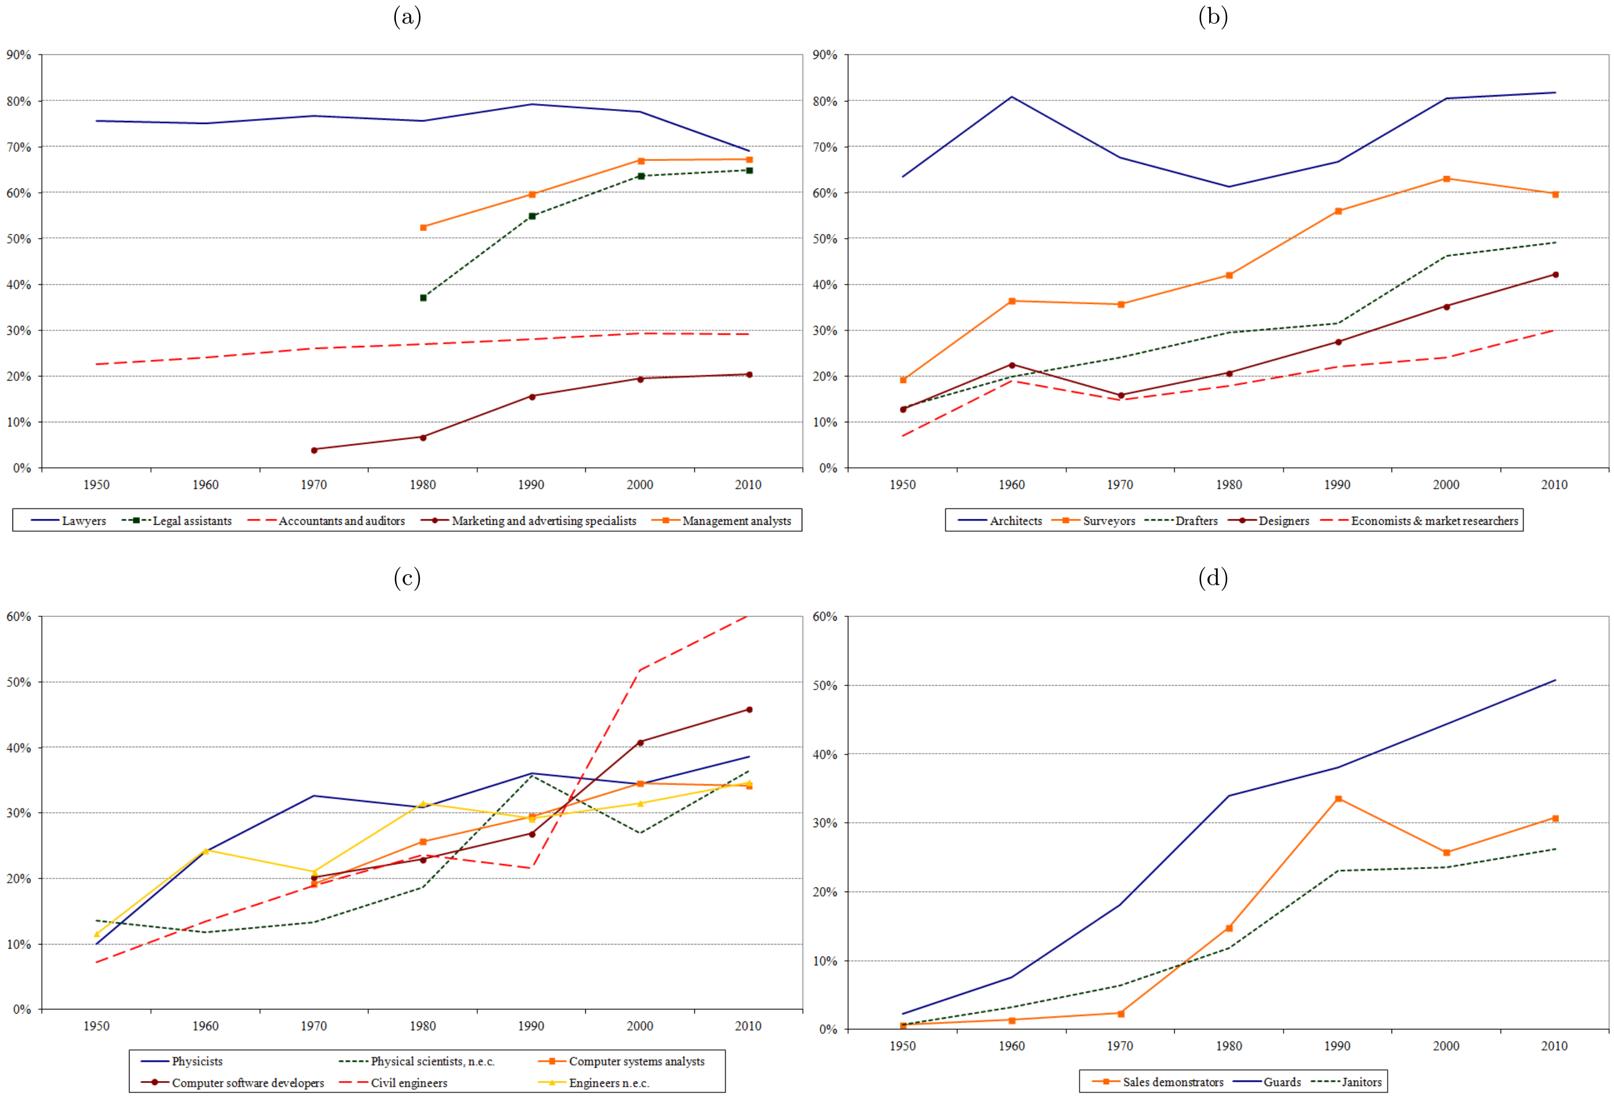

Source: IPUMS-USA.

Appendix

1.A Data

1.A.1 Data Description

1.A.1.1 Industry and I-O Data

All the industry and I-O data come from the Bureau of Economic Analysis (BEA) of the U.S. Department of Commerce. Employment, value added and relative price indexes come from the Annual Industry Accounts, according to the December 2010 release; final uses price indexes come from the National Income and Product Accounts (NIPA) tables. The I-O data for years 1947, 1958, 1963, 1967, 1972, 1977, 1982, 1987, 1992, 1997 and 2002 come from the Benchmark Input-Output Accounts; while data for years 1998-2001 and 2003-2007 come from the Annual Industry Accounts, according to the December 2010 release. Both the standard and the supplementary versions of the tables are considered. The standard versions of the tables are available for years starting from 1992; under this version, the output of industries corresponds to the published output in the Industry Accounts because the redefinitions for secondary products performed by the BEA are not present, as in the supplementary tables. The re-classifications of secondary products carried out by BEA to define commodities cannot be avoided however. I-O tables until 1992 are based on the SIC classification while they are based on NAICS for later years.

The allocation of industries to the three main sectors under investigation is performed as follows:

- Agriculture: Agriculture, forestry, fishing and hunting
- Manufacturing: Mining, Construction, Manufacturing
- Services: all other industries including Government (excluding Scrap, which is kept as a separate sector)

Given the high level of aggregation, the definition of the three main sectors is not heavily affected when the classification switches from SIC to NAICS because most of the changes take place within each aggregate sector. Only two sub-sectors switch from one main sector to another: publishing and auxiliary units. They were both classified within manufacturing under SIC, but are now classified within services under NAICS. Unfortunately it is not possible to perform this adjustment in an ideal way. In particular there is a problem with auxiliary units, which are classified within the sector 55 of NAICS, namely Management of Companies and Enterprises. This sector is composed by three sub-sectors: 551111 (Offices of Bank Holding Companies); and 551112 (Offices of Other Holding Companies); 551114 (Corporate, Subsidiary, and Regional Managing Offices). The latter was moved from manufacturing to PBS but the first two were not. In fact, they were already classified within services under SIC as well. The trouble is that I-O data are not disaggregated enough to distinguish these three sub-sectors, hence, by re-classifying the entire sector within manufacturing, the contribution of PBS is underpredicted.

In the case of publishing the re-classification can be precisely performed by bringing industry 5111 - Newspaper, periodical, book, and directory publishers - back to manufacturing. Yet this can be done for the benchmark years only, because in the case of the Annual I-O Accounts the level of disaggregation is not detailed enough to identify sector 5111; the re-classification has to be performed by moving the entire sector 511 - Publishing Industries (except Internet) - to manufacturing. This latter sector includes 5112 - Software Publishers - that is actually classified in PBS under SIC. This brings about an even more severe underprediction for Annual Accounts, not only for the overall service sector but more importantly for PBS, the main sector of interest in the paper.

The Professional and Business Services (PBS) industry in this study is identified with sector 73 of the SIC I-O classification (until 1992), which includes: 73A (Computer and data processing services); 73B (Legal, engineering, accounting, and related services); 73C (Other business and professional services, except medical); and 73D (Advertising). In terms of the 1987 SIC classification, the sectors included are:

- 73: Business Services:
  - 731: Advertising
  - 732: Consumer Credit Reporting Agencies, Mercantile
  - 733: Mailing, Reproduction, Commercial Art and Photography, and Stenographic Services
  - 734: Services to Dwellings and other Buildings
  - 735: Miscellaneous Equipment Rental and Leasing
  - 736: Personnel Supply Services
  - 737: Computer Programming, Data Processing, and Other Computer Related Services
  - 738: Miscellaneous Business Services
- 76: Miscellaneous Repair Services
  - 769: Miscellaneous Repair Shops and Related Services
- 81: Legal Services
  - 811: Legal Services
- 87: Engineering, Accounting, Research, Management, and Related Services
  - 871: Engineering, Architectural, and Surveying
  - 872: Accounting, Auditing, and Bookkeeping Services
  - 873: Research, Development, and Testing Services (excluding sector 8733 - Noncommercial Research Organizations)
  - 874: Management and Public Relations Services
- 89: Miscellaneous Services

– 899: Miscellaneous Services

The definition of PBS is slightly more restrictive compared to the one employed in the aggregate SIC data presented in Figure 1.4. In particular the following SIC sectors are not included: 762 (Electrical Repair Shops); 763 (Watch, Clock, and Jewelry Repair); 764 (Reupholstery and Furniture Repair); 84 (Museums, Art Galleries, and Botanical and Zoological Gardens); and 8733 (Noncommercial Research Organizations).

The definition of the PBS according to the 2002 NAICS I-O data include sectors: 54 (Professional and Technical Services); 55 (Management of Companies and Enterprises); and 56 (Administrative and Waste Services). The codes coincide with the standard 2002 NAICS codes. This definition does not exactly match the one used under the SIC I-O classification and some adjustments are necessary in order to improve the consistency of the data over time. The re-classification of the sector "Management of Companies and Enterprises" within manufacturing is the first obvious one, given what has just been discussed. Finer adjustments can only be performed for benchmark years because the Annual Accounts lack the needed level of detail; they involve the exclusion of some sub-sectors from the NAICS definition and the inclusion of others that were previously classified within PBS under the SIC definition. Unfortunately it is not possible to get a perfect match; a conservative approach has therefore been used, by moving only sectors whose entire output or the vast majority of it needs to be re-classified. The NAICS I-O sub-sectors that have been excluded from the PBS definition under NAICS are:

- 5615: Travel arrangement and reservation services[41]
- 5620: Waste management and remediation services[42]

The sub-sectors that have been moved to PBS because they belong to it according to SIC are:

- 5112: Software publishers
- 5180: Internet service providers, web search portals, and data processing
- 5324: Commercial and industrial machinery and equipment rental and leasing[43]

Notice that the following SIC sectors cannot be correctly re-classified so they are completely missing from the new definition under NAICS: 7352 (Medical Equipment Rental and Leasing); 7377 (Computer Rental and Leasing); 7378 (Computer Maintenance and Repair); 7383 (News Syndicates); 7384 (Photofinishing Laboratories); and 8741 (Management Services). The vast majority of 769 (Miscellaneous Repair Shops and Related Services) and parts of few other small sub-sectors are missing as well. Instead the NAICS sub-sectors that are kept while they should have been completely dropped because they were not in PBS under SIC are: 541191 (Title

[41] Part of the sector should have been kept because it corresponds to SIC sector 7389 (Business Services, NEC)

[42] Part of the sector should have been kept because it corresponds to SIC sectors 7359 (Equipment Rental and Leasing, NEC) and 7699 (Repair Shops and Related Services, NEC)

[43] This also includes SIC sector 4741 (Rental of Railroad Cars), which was not in PBS; however, the vast majority of it corresponds to SIC sector 735 (Miscellaneous Equipment Rental and Leasing), which is in PBS.

Abstract and Settlement Offices); 541213(Tax Preparation Services); 541921 (Photography Studios, Portrait); 561730 (Landscaping Services); and 561740 (Carpet and Upholstery Cleaning Services).

1.A.1.2 Occupational Data

Occupational data come from the IPUMS-USA database. In order to compare occupations over time, the classification proposed by Meyer and Osborne (2005) is used.[44] The occupations associated with PBS are selected according to the different definitions described in the main text using data in 1990. The list of occupation selected according to the 9% definition are listed in Table 1.A.1. The table also shows the codes corresponding to the categories used to subdivide the occupations. They are:

- 1: Managers
  - 11: Top Managers
  - 12: Other managers
  - 13: Financial Managers
- 2: Professionals
  - 21: Lawyers
  - 22: Architects
  - 23: Engineers
  - 24: Accountants
  - 25: Advertisers
  - 26: Other professions
- 3: Computer related occupations
  - 30: Computer system analysts, software developers etc.
- 4: Clerks
  - 41: Administrative related occupations
  - 42: Service occupations
  - 43: Sales occupations
- 5: Technicians
  - 50: Technicians and repairers
- 6: Other occupations
  - 61: Construction and precision production occupations
  - 62: Operators and laborers

[44] The corresponding variable is named OCC1990.

Table 1.A.1: PBS Occupations - 9% Definition

Occupation Description	**OCC1990**	**Category**
Human resources and labor relations managers	8	11
Managers and specialists in marketing, advertising, and public relations	13	25
Managers and administrators, n.e.c.	22	12
Accountants and auditors	23	24
Management analysts	26	12
Personnel, HR, training, and labor relations specialists	27	12
Business and promotion agents	34	12
Management support occupations	37	12
Architects	43	22
Civil engineers	53	23
Electrical engineer	55	23
Not-elsewhere-classified engineers	59	23
Computer systems analysts and computer scientists	64	30
Operations and systems researchers and analysts	65	30
Statisticians	67	26
Mathematicians and mathematical scientists	68	26
Physicists and astronomers	69	26
Chemists	73	26
Atmospheric and space scientists	74	26
Geologists	75	26
Physical scientists, n.e.c.	76	26
Agricultural and food scientists	77	26
Biological scientists	78	26
Medical scientists	83	26
Economists, market researchers, and survey researchers	166	26
Sociologists	168	26
Social scientists, n.e.c.	169	26
Urban and regional planners	173	26
Lawyers	178	21
Writers and authors	183	26
Technical writers	184	26
Designers	185	26
Art makers: painters, sculptors, craft-artists, and print-makers	188	26
Photographers	189	26
Art/entertainment performers and related	194	26
Editors and reporters	195	26
Electrical and electronic (engineering) technicians	213	50
Engineering technicians, n.e.c.	214	50
Mechanical engineering technicians	215	50
Drafters	217	50
Surveyors, cartographers, mapping scientists and technicians	218	50
Other science technicians	225	50
Computer software developers	229	30

Legal assistants, paralegals, legal support, etc	234	21
Technicians, n.e.c.	235	50
Advertising and related sales jobs	256	25
Sales demonstrators / promoters / models	283	43
Computer and peripheral equipment operators	308	30
Secretaries	313	41
Stenographers	314	41
Typists	315	41
Interviewers, enumerators, and surveyors	316	41
Receptionists	319	41
Information clerks, n.e.c.	323	41
File clerks	335	41
Bookkeepers and accounting and auditing clerks	337	24
Billing clerks and related financial records processing	344	24
Duplication machine operators / office machine operators	345	41
Mail and paper handlers	346	41
Office machine operators, n.e.c.	347	41
Other telecom operators	349	41
Mail clerks, outside of post office	356	41
Messengers	357	41
Customer service reps, investigators and adjusters, except insurance	376	41
Bill and account collectors	378	41
General office clerks	379	41
Proofreaders	384	41
Data entry keyers	385	41
Statistical clerks	386	41
Housekeepers, maids, butlers, stewards, and lodging quarters cleaners	405	42
Supervisors of guards	415	42
Guards, watchmen, doorkeepers	426	42
Supervisors of cleaning and building service	448	42
Janitors	453	42
Pest control occupations	455	42
Small engine repairers	509	50
Repairers of data processing equipment	525	50
Repairers of household appliances and power tools	526	50
Precision makers, repairers, and smiths	535	50
Locksmiths and safe repairers	536	50
Office machine repairers and mechanics	538	50
Mechanics and repairers, n.e.c.	549	50
Paperhangers	583	61
Precision grinders and filers	644	61
Furniture and wood finishers	658	61
Upholsterers	668	61
Photographic process workers	774	62
Welders and metal cutters	783	62
Hand painting, coating, and decorating occupations	789	62

1.A.2 Construction of Aggregate I-O Tables

For the purpose of this study, I-O tables have to be aggregated in order to obtain the I-O linkages for the three main sectors: agriculture, manufacturing and services. The matrix Ω in the model corresponds to an industry-by-industry total requirements table. The methodology to obtain this matrix is described by Horowitz and Planting (2006). In brief, there are two main methods to obtain the matrix corresponding to the different I-O conventions used before and after 1972. For the benchmark years until 1967, one symmetric industry-by-industry transaction matrix is published under the assumption that each industry only produces one commodity and that each commodity is only produced by one industry. The total requirements table is then simply obtained as a Leontief inverse. Since 1972 instead, the symmetry assumption has been dropped and two distinct tables have been published: the commodity-by-industry use table that shows the uses of commodities by industries and final consumers; and the industry-by-commodity make table that shows the production of commodities by industries. The methodology is slightly more involved, but again it is possible to obtain an industry-by-industry total requirements table. In this study, transaction, make and use tables are first aggregated and then inverted to obtain the total requirements table according to the two different methodologies. Moreover, following the documentation for benchmark years, the Commodity Credit Corporation adjustment is performed for years between 1963 and 1977; and the Scrap adjustment is carried out for years between 1972 and 1997.

1.A.3 Construction of the Price Indexes

The aggregated value-added price indexes for agriculture, manufacturing and services have been computed from the chain-type price indexes for value added at the industry level, following the methodology described by Whelan (2002). The price index for agriculture is readily available and corresponds to the aggregate industry "agriculture, forestry, fishing, and hunting". Manufacturing includes the industries "mining", "construction" and "manufacturing". Services include "private services-producing industries" and "government".

The procedure to obtain the final uses price indexes is a bit more involved. All data come from the NIPA tables and since all price indexes are chained, any manipulation described here requires the methodology for chain-type indexes. The procedure involves three main steps: 1) identify the NIPA categories that better represent the I-O definition of commodities; 2) remove transportation, retail and wholesale margins to obtain producers' price indexes; 3) add investment to the relevant sectors and obtain an aggregate price index for each sector that reflects the price of investment as well. The first two steps are described here, while the adjustment for investment is analyzed in Appendix 1.B.3. The first step consists in matching the personal consumption expenditures from the I-O side to the appropriate NIPA categories. Since the NIPA tables were extensively revised in 2009 to incorporate the results of the 2002 benchmark I-O accounts, I perform the match using the 2002 Bridge Table, which links the two data sources. As pointed out in the main text, the identification of agriculture with the NIPA category "food and beverages purchased for off-premises consumption" is not correct because it is seven times larger than personal consumption expenditures for the I-O commodity agriculture; a finer definition is therefore needed. This is achieved by using the underlying NIPA tables, which contain categories at a more disaggregated level. The trouble is that the underlying tables are only available since 1959, hence it is not possible to keep the same exact definition for the three main sectors throughout the entire time period. After 1959, the personal and government consumption expenditures categories are allocated to the three main I-O commodities as follows:

- Agriculture: "Fish and seafood"; "Eggs"; "Fresh fruits and vegetables"; "Food produced and consumed on farms"; "Flowers, seeds, and potted plants"
- Manufacturing: "Durable goods" except "Net purchases of used motor vehicles", "Recording media", "Computer software and accessories" and "Corrective eyeglasses and contact lenses"; "Nondurable goods" except categories already included in Agriculture and "Net expenditures abroad by U.S. residents"; "Food furnished to employees (incl. military)"
- Services: "Services" except "Food furnished to employees (incl. military)"; "Recording media"; "Computer software and accessories"; "Corrective eyeglasses and contact lenses"; "Net expenditures abroad by U.S. residents"; "Government consumption expenditures"[45]

The match cannot be perfect because each NIPA category is often associated with more than one I-O commodity. For instance, "Cereals" are allocated in part to "Crop products", which fall in agriculture, and in part to "Food products", which fall in manufacturing. A conservative approach is used and a category is moved only if the majority of its expenditures falls in another sector. In the case of "Cereals", they are moved to manufacturing because only 1% of their expenditures are associated to agricultural commodities. Despite the imperfect match, the magnitudes are now much more in line with I-O data; for instance the personal consumption expenditures allocated to agriculture amount to 47.4 billions of dollars (at producers' prices) in 2002 while they are 48.2 billions of dollars in the I-O data. Unfortunately the same level of disaggregation is not available before 1959 and a much coarser match has to be used.[46] The three main sectors are identified as follows:

- Agriculture: "Food and beverages purchased for off-premises consumption" except "Alcoholic beverages purchased for off-premises consumption"
- Manufacturing: "Durable goods" except "Net purchases of used motor vehicles"; "Nondurable goods" except categories already included in Agriculture; "Food furnished to employees (incl. military)"

[45] The treatment of government consumption expenditures changed in 1998. The reason is that the gross output for the general government industry did not include intermediate inputs before 1998 and they were accounted for as government consumption expenditures. Therefore the complete association of government consumption expenditures with services is correct only in recent years. Before 1998, one should allocate part of the government expenditures to agriculture and manufacturing; unfortunately the Bridge Tables are not available for government consumption expenditures and it is not clear which NIPA categories should be reallocated. In any case this is unlikely to have a major impact; in fact the government expenditures on agriculture were almost nil in all years and the expenditures on manufacturing commodities that should be reallocated were just 15% of the total in 1997.

[46] As a robustness exercise, in order to exclude this initial period, the main results of the paper are replicated starting from the benchmark table in 1958. They are very robust if not stronger. In fact PBS outsourcing accounts for 2.4 percentage points of the change; given the shorter period this corresponds to 14.4% of the total increase in the share of services in total employment.

- Services: “Services” except “Food furnished to employees (incl. military)”; “Government consumption expenditures”

The price indexes obtained so far are in purchasers' prices, however; this implies that part of their value reflects margins that actually belong to the service sector. The second step therefore consists in obtaining the transportation, retail and wholesale margins for agriculture and manufacturing from I-O tables. The data are available only for benchmark years starting from 1967; thus interpolated values are used in missing years and the margins for the 1947-1966 period are assumed to be equal to their value in 1967. The agriculture and manufacturing price indexes are adjusted to remove these margins, which are then moved within services. To achieve this, price indexes for transportation, retail and wholesale trade are needed. For transportation I take the price index for “Public Transportation” from NIPA tables. For retail and wholesale trade instead there is no direct counterpart in the NIPA tables (there is no final demand for retail trade as such). The obvious choice would be to take price indexes for gross output from the Industry Accounts; unfortunately gross output prices are available only since 1987, therefore valued added price indexes are used instead.

1.B Extra Results

1.B.1 Results until 2007

In recent years, the I-O tables are available annually and not only for the benchmark years. Unfortunately, the annual tables are computed using more aggregate data and do not match the statistical quality of tables in benchmark years. In particular, the intermediate inputs at the detail level are estimated assuming the industry technology to be constant, undermining the precise aim of this study. Moreover, the annual tables are revised periodically over time[47], when new information becomes available; instead the benchmark tables are usually published with a 5-year lag and are not subject to further updates. Also the correction for the classification change cannot be performed as precisely as for benchmark years, as pointed out in Appendix 1.A.1.1. The finer adjustment for PBS cannot be done; and, in the case of publishing, I have to re-classify a larger sector that includes Software Publishers, causing an even bigger underprediction of the overall service sector. For all these reasons, the data for years after 2002 are particularly inaccurate, and the results should be therefore treated with care.

Table 1.B.1: Predicted versus Actual Changes in Employment Shares - No Auxiliaries

Sector	Data	Prediction	Ratio
Agriculture	-4.05	-3.27	81%
Manufacturing	-19.35	-3.15	16%
Services	23.41	6.41	27%

Note: Period: 1948-2007. See also notes in Table 1.3.

[47] This study uses data from the December 2010 revision.

Table 1.B.2: Effect of Outsourcing on the Service Employment Share - No Auxiliaries

Counterfactual	Predicted Change	Ratio to Baseline	Diff. wrt Baseline
Baseline	6.41	100%	0.00
1: No Service Outsourcing	1.90	30%	4.51
2: No PBS Outsourcing	3.87	60%	2.55
3: No Finance Outsourcing	6.13	96%	0.28

Note: Period: 1948-2007. See also notes in Table 1.4.

I replicate the results of Section 1.4.2.3 over the period 1948-2007. As expected, given the warning on data quality, the predictions drop slightly in recent years. As shown in Table 1.B.1, the predicted change in the service share is equal to 6.41 percentage points of total employment, which corresponds to 27% of the actual change. An extra reason for the drop in the estimate is that, after having somewhat leveled in the '90s, the employment share of services experienced a sharp increase in the last decade. Looking at the contribution of outsourcing in Table 1.B.2, PBS still account for around 40% of the total. The contribution is lower in absolute terms, 2.55 percentage points of total employment, but is not far from the 3 percentage point change obtained in the main results. Despite the data quality issues, PBS outsourcing still accounts for a sizable share of the total labor reallocation.

1.B.2 Results with Standard I-O Tables

This appendix shows the results obtained using the standard I-O tables. In these tables output of industries corresponds to the published output in the Industry Accounts because the redefinitions for secondary products performed by the BEA are not present. As a robustness exercise, I report the estimates obtained using these tables for the change in the employment share until 2002. Tables 1.B.3 and 1.B.4 show the results of the exercise, which is performed according to the setting of Section 1.4.2.3 where the elasticity was fixed to one in order to isolate the forces under study. Tables that replicate results of other sections of the paper are available on request; they are not reported here because they do not add any extra evidence. As expected, there is almost no impact on the contribution of outsourcing; here PBS outsourcing accounts for 40% of the change, against 41% in the results reported in the main text. The impact of the redefinition is mainly on the magnitude of the results, but again it is very marginal. The proposed gross output model is capable of explaining a change in the service share equal to 7.3 percentage points of total employment in 2002, versus the 7.4 percentage points found when supplementary tables are used. In absolute terms, outsourcing accounts for 2.9 percentage points, just 0.1 percentage points less than before.

Ideally one would like to obtain the results using tables that exactly match the Industry Accounts data, that is, tables without the re-classification of secondary products performed by BEA to define commodities. These re-classifications are the main reason why it is not possible to affirm as strongly as for industry data that the change in PBS coincide with a rise in outsourcing.

Table 1.B.3: Predicted versus Actual Changes in Employment Shares - Standard Tables - No Auxiliaries

Sector	Data	Prediction	Ratio
Agriculture	-3.99	-3.45	86%
Manufacturing	-18.28	-3.85	21%
Services	22.28	7.30	33%

Note: Period: 1948-2002. See also notes in Table 1.3.

Table 1.B.4: Effect of Outsourcing on the Service Employment Share - Standard Tables - No Auxiliaries

Counterfactual	Predicted Change	Ratio to Baseline	Diff. wrt Baseline
Baseline Model	7.30	100%	0.00
1: No Service Outsourcing	3.01	41%	4.28
2: No PBS Outsourcing	4.41	60%	2.89
3: No Finance Outsourcing	7.10	97%	0.20

Note: Period: 1948-2002. See also notes in Table 1.4.

Unfortunately tables before re-classifications are not published. However, as pointed out in Section 1.2.2, these re-classifications mainly affect small single-establishment firms and internal transactions seem to be constant over time, so they are unlikely to have a strong impact on the change. The results obtained for tables before redefinitions offer further strong evidence. In fact, the redefinitions are performed using exactly the same logic of the re-classifications, only they are applied to the definition of industries and not commodities. The very small impact of these redefinitions on the magnitude of the results is reassuring and proves that what is observed in the data is mainly driven by outsourcing: similarly the re-classifications are likely to have a very marginal impact.

1.B.3 Results with Investment

Despite being by far the largest component (85.6% in 2002), personal and government consumption expenditures do not account for the total of final uses, and a further extra adjustment in the price indexes is needed in case investment is to be considered as well. This adjustment involves the allocation of private fixed investment and government gross investment to the three main sectors. The agriculture sector is not a recipient of investment, so no further modification is needed. Unfortunately the NIPA tables are again not detailed enough, and the allocation is quite coarse. All of investment apart from investment in software is allocated to manufacturing; hence the investment allocated to services are just software plus the transportation, retail and wholesale margins associated with investment in manufacturing.[48] The share of investment allocated to services is therefore lower than the actual one. For instance, part of the investment

[48]Margins for fixed private investment and government gross investment are again obtained from benchmark I-O tables and interpolated in missing years. Unfortunately the first year in which these margins are available is 1982; hence in all previous years the margins are assumed to be equal to their value in 1982. This does not seem to be a particular source of concern given that the margins are quite constant over time.

in structures should be allocated to Real Estate, which is in services; PBS is another recipient of investment, which cannot be clearly identified.

The results of the paper are re-obtained here to show the robustness to the inclusion of investment. Clearly the relevant results are those of Section 1.5 when the final uses expenditure shares are allowed to vary, since with a unitary elasticity the choice of the price indexes do not matter at all. An extra step is required to allow for investment in the value added model, otherwise the comparison between the two models would not be correct. The treatment of investment in the value added model is performed as in Ngai and Pissarides (2004); they assume that all of investment is performed in manufacturing and set the aggregate investment rate to 20% of output, matching the average investment rate for the period 1929-1998. Note that this is similar to the adjustment performed for the gross output prices, since, also in that case, the share of investment performed in the service sector cannot be properly accounted for.

Table 1.B.5: Predicted vs. Actual Changes in Employment Shares - Investment and No Auxiliaries

		Gross Output		Value Added	
Sector	Data	Prediction	Ratio	Prediction	Ratio
Agriculture	-3.99	-3.53	88%	-2.90	73%
Manufacturing	-18.28	-8.01	44%	-3.01	16%
Services	22.28	11.54	52%	5.90	27%

Note: The predicted changes are obtained using both the proposed Gross Output model and the Value Added benchmark model. Period: 1948-2002. The elasticity of substitution $\epsilon = 0.5$. See also notes in Table 1.3.

Table 1.B.6: Effect of Outsourcing on the Service Employment Share - Investment and No Auxiliaries

	Predicted Change		Extra Prediction		
Counterfactual	Value Added	Gross Output	Difference	Ratio to Baseline	Diff. wrt Baseline
Baseline	5.90	11.54	5.64	100%	0.00
1: No Service Outsourcing	5.90	7.52	1.62	29%	4.03
2: No PBS Outsourcing	5.90	8.90	2.99	53%	2.65
3: No Finance Outsourcing	5.90	11.36	5.46	97%	0.18

Note: The Extra Prediction is defined as the difference between the employment share change predicted by the proposed Gross Output model and the change predicted by the Value Added benchmark model. Period: 1948-2002. The elasticity of substitution $\epsilon = 0.5$. See also notes in Table 1.4.

Tables 1.B.5 and 1.B.6 report the results of the exercise. The overall predicted sectoral reallocation is reduced in both models; this result comes from the fact that most of the investment is accounted for in manufacturing, hence this sector experiences a lower drop in total employment. In fact, according to the gross output model, the change in the share of manufacturing is equal to -8.01 percentage points of total employment in 2002, a lower drop compared to the 9.42 points predicted in the main text without accounting for investment. Also the predicted increase in services is lower, amounting to 11.54 percentage points versus the 12.98 points predicted without investment. But the contributions of the change in the I-O structure and of outsourcing are very robust, displaying even higher values compared to the results without investment. In fact, accounting for intermediates improves the prediction of the rise in the service share by 5.64 percentage points. For what concerns the contribution of outsourcing the results are also robust, if not stronger. Service outsourcing potentially accounts for 71% of the total extra prediction; and if the contribution is more plausibly narrowed to PBS only, outsourcing explains 47% of the total. This is a smaller share compared to the 53% in the main text, but it corresponds to a higher amount in absolute terms: 2.65 percentage points of total employment compared to 2.52 points predicted without including investment.

1.C Determinants of PBS Outsourcing: Census data

The measure of purchased PBS used in Section 1.6.2 is obtained from I-O tables. As argued in the main text, this measure of PBS outsourcing is reliable once auxiliary units are excluded; in fact, the problem of internal transactions only remains for those small companies whose secondary products are re-classified by the BEA from manufacturing to PBS. These transactions are likely to account for a very small share of the total. In any case, to dispel any doubt on this issue I perform a robustness exercise and use a second more precise measure of service outsourcing. It comes from the quinquennial Census of Manufactures, which directly asks firms the cost of services purchased from other companies. The problem of internal transactions is therefore completely eliminated. Unfortunately the first year in which data are available is 1992, and only a limited range of services is available: legal, accounting, advertising, software and data processing, and refuse removal. These constitute a subset of the services contained in the PBS sector.

The industry classification employed is NAICS, and I convert the data in 1992 from SIC to NAICS using the weighted concordance table available on the U.S. Census Bureau website. The measure of coordination complexity is obtained using the Occupational Employment Statistics published by the U.S. Bureau of Labor Statistics. The data are available at a 4-digit NAICS level only from 2002, therefore I cannot exploit the within variation and the analysis only focuses on the cross-sectional variation by adding year fixed effects. A further reason for this choice is that the measure of service outsourcing is not completely consistent across the different Censuses; in fact the 2002 Census also includes purchases of computer hardware, which cannot be excluded[49].

Table 1.C.1 shows the results of the regressions. Coordination complexity again has a strongly positive and significant effect on PBS outsourcing. The adoption of new technologies, measured by the number of patents used by the industry, has a positive effect but not robust to the inclusion of all controls. Allowing for cross-industry variation only, I can include other determinants of outsourcing, whose measure is only available in a given year. They include: a measure of productivity dispersion as in Yeaple (2006); the ratio of R&D expenditures

[49]Data in 2002 also include the cost for management consulting and administrative services. Since the time variation is not exploited, they are not excluded because they are contained in PBS.

Table 1.C.1: Determinants of PBS Outsourcing - Census data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Complexity	1.909^a	1.478^b	1.564^b	3.554^a	2.590^a	2.426^a	2.386^a	2.357^a	2.538^a	2.783^a
	(0.544)	(0.627)	(0.636)	(0.484)	(0.475)	(0.468)	(0.475)	(0.474)	(0.474)	(0.497)
Num Patents		0.071^a	0.062^a	0.090^a	0.049^a	0.053^a	0.043^b	0.035^c	0.028	0.031
		(0.021)	(0.021)	(0.018)	(0.019)	(0.019)	(0.020)	(0.020)	(0.020)	(0.020)
Num Inputs			0.151	0.229^b	0.201^b	0.219^b	0.223^b	0.236^b	0.194^c	0.217^b
			(0.113)	(0.097)	(0.099)	(0.102)	(0.102)	(0.103)	(0.105)	(0.105)
K/L				-0.406^a	-0.406^a	-0.330^a	-0.322^a	-0.308^a	-0.245^a	-0.241^a
				(0.031)	(0.031)	(0.057)	(0.059)	(0.063)	(0.067)	(0.067)
S/L					0.302^a	0.290^a	0.261^a	0.277^a	0.254^a	0.280^a
					(0.049)	(0.048)	(0.054)	(0.055)	(0.056)	(0.060)
Scale						-0.052	-0.058	-0.074^c	-0.079^c	-0.076^c
						(0.034)	(0.036)	(0.042)	(0.042)	(0.042)
R&D/Sales							0.038	0.042	0.024	0.035
							(0.027)	(0.027)	(0.027)	(0.029)
Dispersion								0.054	0.058	0.051
								(0.057)	(0.055)	(0.055)
Contract Int									0.151^a	0.148^a
									(0.053)	(0.053)
Routine										0.407
										(0.302)
Observations	1,386	1,383	1,383	1,376	1,376	1,376	1,367	1,352	1,352	1,352
R-squared	0.043	0.062	0.064	0.229	0.263	0.265	0.268	0.279	0.286	0.287
Fixed effects	year	year	year	year	year	year	year	year	year	year

Note: The dependent variable is the share of purchased professional and business services from other companies over total sales. All variables are expressed in logs. Data are from the Census of Manufactures for years 1992, 1997 and 2002. Industry-clustered standard errors in parentheses; (a, b, c) indicate 1, 5, and 10 percent significance levels.

to sales from the FTC Line of Business Survey; the measure of contract intensity proposed by Nunn (2007); and the measure of routine introduced by Costinot et al. (2011). Analyzing the control variables, human-capital intensity again has a positive effect, and this time it is strongly significant. Capital intensity is instead negative and significant, in contrast with the previous results that gave a positive estimate. The positive and significant effect of the contract intensity variable can be interpreted as another support, albeit indirect, to the complexity and core-competencies story. Under the standard Property Right Theory interpretation, a firm will in-source more contract intensive inputs. Given that all of the inputs used to construct this variable are goods, the positive impact on service outsourcing can be rationalized by arguing that a manufacturing firm with more contract intensive inputs will focus on its core-competencies by producing more goods in-house and outsourcing more of the non-core services.

1.D The Rise of Business Services in OECD Countries

In this Appendix I show that the role played by business services in the rise of the service sector is roughly the same, if not larger, across several OECD countries. Business services are the industry that accounts for the largest share of the total rise in services, both in terms of GDP and employment. In the average country they account for 53% of the increase in services in terms of percentage points of GDP, with this share rising to more than 75% for Canada and Denmark (Table 1.D.1). Considering the rise in terms of total employment (Table 1.D.2), business services account for 44% on average. It is interesting to notice that the contribution drops very little compared to Finance and Real Estate; in the Netherlands it even increases to 69% of the total.

Combining these findings with the fact that in all countries more than 90% of the output of business services constitutes an intermediate product to firms, I find that input-output linkages evolve very similarly in all countries of the sample. Figure 1.D.1 shows the forward linkage for business services and the median industry.[50] This is a measure of the interconnection of a sector to all other sectors through the supply of intermediate inputs. In light of the insights offered by Acemoglu et al. (2012), a sector with a larger forward linkage has a stronger influence on the rest of the economy since any shock to it will propagate to a large part of the economy. The quantity plotted in Figure 1.D.1 essentially corresponds to an element of what Acemoglu et al. (2012) define as "influence vector" (up to the labor share). Business services have experienced a remarkable growth in their forward linkage, becoming the industry with the highest influence on the rest of the economy in all countries of the sample.

This transformation might have important implications for aggregate outcomes. First of all any productivity gain accrued in this sector will have large spillovers on the downstream sectors, affecting aggregate productivity. Intuitively, once a gross output setting is adopted, the contribution of each sector to aggregate productivity is not simply based on valued added shares but will consider how strongly each sector affect the rest of the economy through input-output linkages. If sectoral linkages are constant over time there is essentially no difference in the decomposition of aggregate productivity growth, but, in the presence of some variation, sectors might have different weights once sectoral linkages are considered. In an on-going research project, I find that business services play a much more important role than originally thought, and they explain a large share of the productivity differences that we observe across advanced economies.[51] In essence business services can constitute either a drag or a boost for the rest of the economy depending on their productivity growth.

[50] More precisely the figure plots the forward linkage divided by the total number of industries. The forward linkage is the horizontal sum of the coefficients in the total requirements table or Leontief inverse, which shows the inputs required, both directly and indirectly, from all industries in the economy to produce a dollar of output of each commodity.

[51] See Berlingieri (2013).

Table 1.D.1: The Rise in Services in terms of GDP

	Total Services	Finance		Real Estate		Business Services	
Country	Change	Change	Share	Change	Share	Change	Share
Australia	14.63	3.84	26%	5.66	39%	5.77	39%
Canada	4.41	2.42	55%	0.43	10%	3.36	76%
Denmark	6.58	0.59	9%	2.72	41%	5.84	89%
France	18.57	0.78	4%	6.02	32%	5.64	30%
Germany	15.77	0.62	4%	6.13	39%	9.35	59%
Italy	15.36	0.82	5%	5.80	38%	6.13	40%
Japan	15.75	1.78	11%	3.92	25%	5.91	38%
Netherlands	13.65	2.70	20%	4.23	31%	8.43	62%
United Kingdom	23.10	4.34	19%	4.86	21%	9.67	42%
Average			*17%*		*31%*		*53%*

Note: The change is expressed in percentage points of GDP. The time period if 1970-2007 apart from Canada (1970-2004) and Japan (1973-2006). Data from EU KLEMS.

Table 1.D.2: The Rise in Services in terms of Employment

	Total Services	Finance		Real Estate		Business Services	
Country	Change	Change	Share	Change	Share	Change	Share
Australia	20.75	1.04	5%	0.85	4%	7.07	34%
Denmark	17.51	0.48	3%	0.67	4%	7.37	42%
Canada	11.41	1.04	9%	-0.07	-1%	3.94	34%
France	19.86	0.52	3%	0.56	3%	9.38	47%
Germany	22.45	0.49	2%	0.82	4%	9.74	43%
Italy	18.51	1.32	7%	0.15	1%	7.42	40%
Japan	16.35	-0.53	-3%	0.78	5%	8.27	51%
Netherlands	16.17	0.77	5%	0.48	3%	11.15	69%
United Kingdom	26.14	1.84	7%	1.06	4%	9.64	37%
Average			*4%*		*3%*		*44%*

Note: The change is expressed in percentage points of total employment. The time period if 1970-2007 apart from Canada (1970-2004) and Japan (1970-2006). Data from EU KLEMS.

Figure 1.D.1: Business Services Forward Linkage

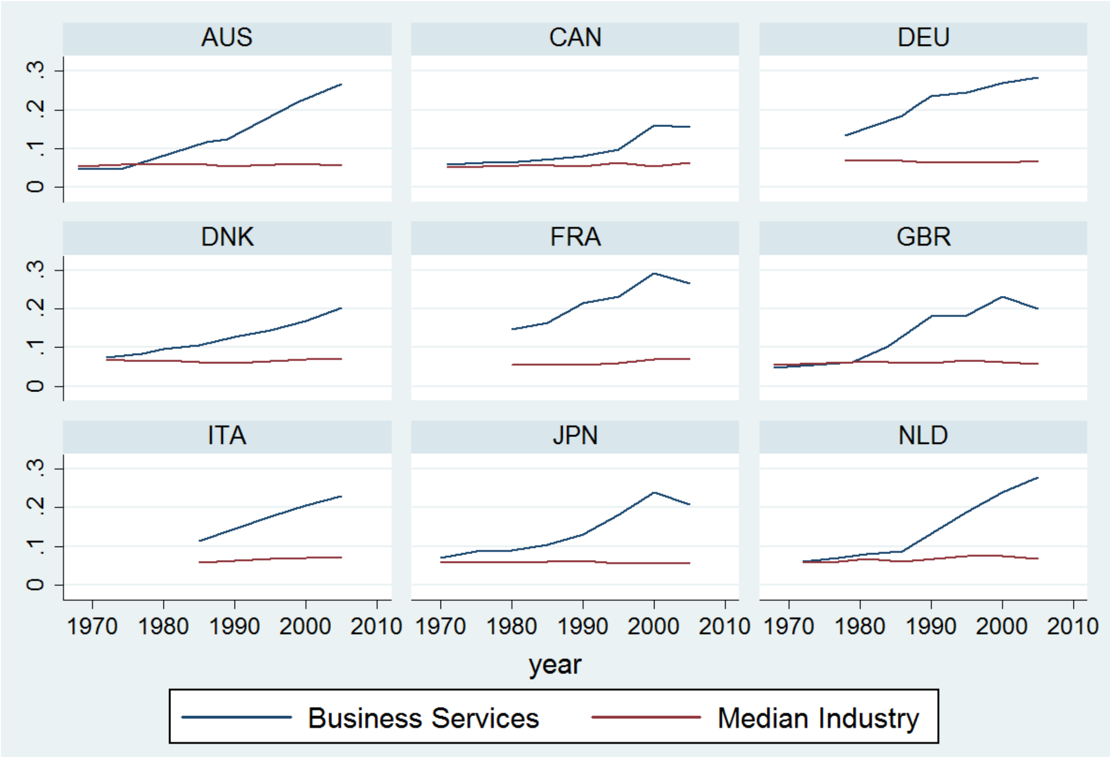

Source: OECD Input-Output database and author's calculations.
Note: The figure plots the forward linkage divided by the total number of industries.

Figure 1.D.2: Business Services Domar Weight

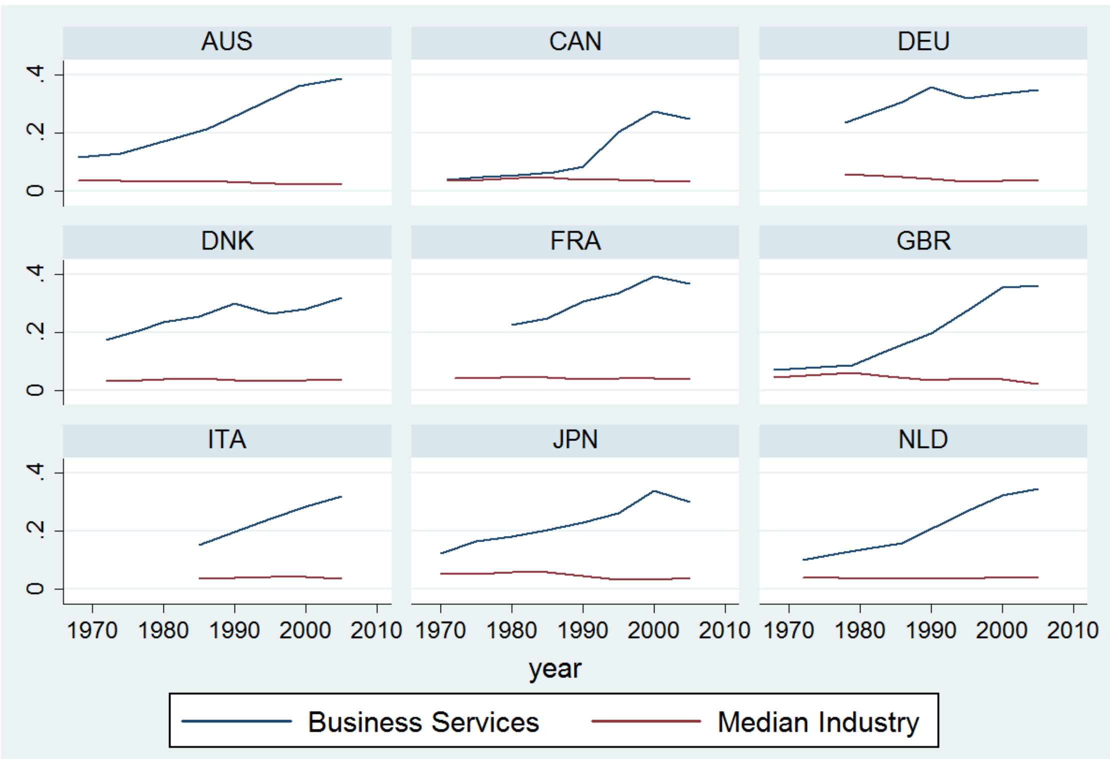

Source: OECD Input-Output database and author's calculations.

Chapter 2

Exporting, Coordination Complexity, and Service Outsourcing

2.1 Introduction

Firms have become more specialized over time. As a consequence, more and more processes and components have been handed over to external specialists, contributing to the growth of outsourcing. Although this is a sensible statement there is no systematic analysis on the trend of domestic outsourcing, as pointed out by Antràs and Helpman (2004). However, a clearer picture emerges when outsourcing is narrowed to the contracting out of services, and business services in particular. Over the past few decades, firms have purchased more and more services from external providers; namely accounting, engineering, legal services but also security, maintenance, janitorial services just to cite few. These services are classified within Professional and Business Services (PBS), and this sector has experienced a dramatic increase. In France, the share of PBS in total GDP was 5.4% in 1970, while the same share was 14.7% in 2007; this almost threefold increase accounts for 47% of the growth of the entire service sector.[1] The pattern is by no means specific to the French economy, a very similar picture holds true for the U.S., the U.K., and many other developed countries.[2] Moreover final demand plays a very marginal role in this rise. The PBS sector is in fact unusual in this regard: in 2005 roughly 94% of its output was used by firms, either as intermediate inputs or in the form of investment, highlighting the primary role played by firms.[3] Understanding what determines the firm's decision to contract out its service inputs is therefore key to explain the causes of the rise of this sector and of services in general.

Despite many studies having focused on service off-shoring, the vast majority of services is actually contracted out domestically.[4] In 2005, business services purchased internationally by French firms accounted for just 7% of the total output of this sector. The small role still played

[1] Data from the EU KLEMS database.

[2] In Chapter 1, I show that this increase in vertical specialization has a sizeable impact on the reallocation of labour across sectors in the U.S., with business services outsourcing alone accounting for 14% of the total increase of the service sector.

[3] Data from the OECD Input-Output database.

[4] On service off-shoring see, among others, Görg et al. (2008), Amiti and Wei (2009), and Jensen and Kletzer (2010).

by international trade in services justifies the focus on the firm boundary dimension. Mainly for data limitations, I do not intend to distinguish between domestic and international outsourcing; but since the vast majority of the service inputs is outsourced domestically, what I observe in the data almost coincides with domestic outsourcing. Moreover most of the literature has focused on the consequences of service outsourcing.[5] With few exceptions (e.g. Abraham and Taylor, 1996), very little attention has been devoted to the determinants of service outsourcing. The goal of this paper is to analyze the key forces that affect the firm's decision to contract out its service inputs, and in doing so I unveil new systematic evidence about domestic service outsourcing using an extensive dataset of French firms. In particular, I find that an increase in the number of export destination countries has a strong positive effect on the share of purchased business services in total costs, even after controlling for total exports and for many other determinants of outsourcing already proposed in the literature. Moreover the causal estimates show that the effect is quantitatively very significant.

In order to rationalize these facts, a model of the boundaries of the firm is needed. The Grossman-Hart-Moore property-rights model, well-established in the trade literature thanks to Antràs (2003), draws the boundary of the firm on the basis of which party owns the asset. But asset ownership is less important in the case of services. Therefore this paper embraces a vision of the firm where the residual rights are mainly in terms of control over the decisions to be taken, and not over the assets. I do so by adopting a Transaction Cost Economics (TCE) and moral hazard view of the firm that stresses the importance of ex-post inefficiencies and of monitoring the actions of the agents. Ex-post adaptation will be at centre stage and the residual rights of control are interpreted as the decision rights to choose the best action in the interest of the organization as a whole.

The contribution of the paper is to incorporate the cost of integration as originally stressed by Ronald Coase. In his celebrated article of 1937, Coase argues that a firm is a method of *coordinating* production that is alternative to the market; and the reason why firms exist is because there are costs associated with using the price coordination mechanism. I adopt coordination complexity as the main ingredient for both market and internal transactions, in the spirit of Becker and Murphy (1992). Ideally all tasks would be coordinated in the market, as the price provides everything "participants need to know to be able to take the right action" (Hayek, 1945, p. 527), and a transaction can be carried out independently from all the others. Unfortunately the way each single input is produced (which can be the most efficient when the input is produced independently) might not fit the overall firm's production process and some adaptation is needed ex-post. It is in this setting that the internal hierarchy overcomes the market: the manager has the ability to steer and coordinate the actions of the employees to implement the best action when adaptation is needed. As in the work of Bajari and Tadelis (2001), the integration decision is driven by the trade-off between the ex-ante price and the ex-post adaptation costs. If the input is purchased from the market, the ex-ante cost will be

[5] For instance, Siegel and Griliches (1992), Fixler and Siegel (1999), ten Raa and Wolff (2001), using industry level data for the U.S., find that TFP growth in the manufacturing sector is positively related to service outsourcing.

low thanks to the high-powered incentives, but the ex-post cost that has to be sustained when adaptation is needed will be very high. On the other hand, producing in-house by employing the supplier reaches precisely the opposite result: the cost will be high because the employee has to be compensated for taking an action that is not ideal for his own task, but the extra ex-post cost when intervention is needed will be low, thanks to better coordination and authority.

Then why is not all production carried on by one big firm? TCE rules out the possibility of a firm growing indefinitely assuming that selective intervention is severely limited: a firm cannot simply outsource the production of tasks ex-ante to capture the benefits of higher incentives and then internalize the modifications in case adaptation is needed. I propose an extra reason based on the limits that bounded rationality imposes on the managerial ability to coordinate production. As noticed by Winter (1988), bounded rationality is at the heart of TCE. But the TCE literature has mainly appealed to bounded rationality to justify the existence of contract incompleteness.[6] In this paper I adopt bounded rationality to highlight the limits of coordination following Crémer et al. (2007). Even allowing for the possibility of selective intervention, the action of the manager still suffers from diminishing returns. Intuitively, if the manager has to coordinate more tasks, she will inevitably become less effective in carrying out the needed adaptation, and the cost of internal production will rise.

The literature has so far analyzed transactions independently, “a series of separable make-or-buy decisions”, as pointed out by Williamson in his Nobel Prize Lecture.[7] In the present setting tasks will be interdependent: the inclusion of a new task hinder the performance of others. Inside the firm, the language is the coordinator device and the manager will choose the optimal code to deal with the problems she faces. Adding a new task implies that the words used to communicate (which are limited in number given bounded rationality) will have to be more generic, making it harder to diagnose all other tasks identified by the same word.[8] Therefore integration costs depend on the number and type of activities already produced by the firm. This brings about the definition of coordination complexity put forward in this paper: the higher the number of tasks that the manager has to supervise, the lower the frequency of each of them and hence the higher the complexity of the environment. In this respect, integration costs decrease when the firm reduces the number of tasks internally produced.

I propose one possible driver of coordination complexity: the internationalization decision of the firm. And I will mainly, but not exclusively, focus on the the service inputs that a manufacturing firm needs to produce its products. The main reason for this choice is that fixed export costs are often characterized as the service inputs needed to export to a particular country; hence exporting to more destination countries implies that more inputs are needed (e.g.: a different advertising campaign for each destination market).[9] Each of these country-

[6]Yet, in an insightful early paper, Williamson (1967) resorts to the bounded rationality to build a hierarchical model of the firm where the size is limited due to the managerial loss of control.

[7]“Transaction Cost Economics: The Natural Progression”, 2009.

[8]Note that the design of a common language does not only capture the mere communication costs of passing a message but also the larger cognitive costs of interpreting and understanding that message; hence there is a tight relationship with the cognitive skills a manager is endowed with.

[9]In motivating the presence of some fixed costs to exporting, Melitz (2003) asserts that a firm must inform foreign buyers about its product, learn about the foreign market, research the foreign regulatory environment etc...

specific service inputs is a low probability event from the point of view of the manager of the manufacturing firm; and if a firm exports to more countries the probability of each event will decrease, which translates into a more complex business environment. The model will then predict that the share of outsourced inputs in total costs increases because coordinating these infrequent tasks in-house would require a very costly communication code. Therefore I proxy the firm's coordination complexity with the number of export destination countries. This choice is very much in line with the most common definition of complexity in systems theory, where complexity arises through connectivity and the inter-relationships of a system's constituent elements.

I empirically test the model using a panel of French firms over the period 1996-2007 and a more detailed survey on service outsourcing available in 2005. Over the entire period I observe purchases of selected business service inputs from other firms, like purchases of studies, IT services, advertisement etc...; while in 2005 I can observe 35 specific types of service inputs. I find that coordination complexity, measured as the number of export destination countries, has a strongly positive and significant effect on the share of purchased business services. The result holds on both the cross-sectional and the within-firm variation, and it is extremely robust to internal production and to the inclusion of alternative determinants of outsourcing proposed in the literature, including size, and capital, skill and contract intensities. I also find that outsourcing of services is not driven by the trade intensive margin, so I provide direct evidence for the widespread assumption that service inputs are a fixed export cost component. I contribute to opening the black box of fixed export costs by showing the precise service inputs a firm needs when exporting, and show that firms tend to acquire these key inputs by outsourcing them to external providers, rather then producing them in-house. I also shed some light on the nature of these costs, showing that they are primarily sunk, rather than fixed costs incurred each period.

Moreover, drawing on the insights of the multi-product literature and assuming product specific fixed export cost, I find that an increase in the number of exported products as well as its interaction with the number of destination countries lead to a higher share of outsourced services.[10] I also find the same overall results when I analyze the outsourcing of non-core activities. The model does not differentiate the inputs; therefore there is no 'a priori' clear distinction between a service or a non-service task, apart from the intuitive assumption that, for manufacturing firms, the importance of adaptation will be higher for the primary good inputs. Coordination complexity has again a positive and even stronger impact on outsourcing, showing that the results generalize to other types of inputs as well.

Finally, I investigate the causal effect of globalization on structural transformation through the outsourcing of business services. I propose a set of firm-level instruments that exploit the information on the product space of the firm and rely on plausibly exogenous demand shocks

These tasks correspond to advertising, market and legal research, and they are all supplied by the professional and business industry. Das et al. (2007) and Morales et al. (2011) put forward very similar arguments. Among others, Eaton et al. (2011) and Helpman et al. (2008) adopt settings that feature country-specific fixed export costs.

[10] Bernard et al. (2011) argue that product-specific fixed costs capture the market research, advertising, and regulation costs that need to be incurred when exporting a product.

as shifters. The new channel I put forward is not only present but it is also quantitatively very significant. The causal effect of globalization essentially explains almost all of the increase in business service outsourcing observed in the sample. This new channel also offers new supporting evidence for the 'learning-by-exporting' hypothesis in the trade literature (e.g. De Loecker, 2007): an extra reason for why firms might experience an increase in productivity after entering new markets is that these firms outsource their service inputs to more efficient external providers.

This paper is related to the recent literature on firm organization and vertical hierarchies (e.g.: Garicano, 2000; Garicano and Rossi-Hansberg, 2006; Caliendo and Rossi-Hansberg, 2012). I assume a very simple type of hierarchy: only two layers with a manager who directs and coordinates her employees. Instead of analyzing the vertical dimension of the firm, I look at the horizontal one and take the boundary of the firm explicitly into account to investigate whether a task is produced internally or outsourced. Since those papers do not explicitly draw the boundary of the firm, there is nothing that imposes that problem solvers, who have the knowledge to solve exceptional problems, should be employed directly by the firm.[11] By taking the horizontal dimension explicitly into account, I can explain why Caliendo et al. (2012) do not find empirical support for all theoretical predictions of the model (e.g. rate of expansion of higher layers), and show how outsourcing allows firms to be more flexible, smoothing the transition between different number of layers.

The paper is organized as follows. In the next section, I provide evidence for the aggregate trend in service outsourcing observed in recent years. Section 2.3 reviews some of the key contributions in the literature of service outsourcing, while the following section presents the model. In Section 2.5, I test the main predictions of the model using firm-level data from France and in the following section I present some extra interesting predictions on the non-linear behavior of the outsourcing share. Section 2.7 concludes; some extensions to the baseline model, the description of the data and some extra results are contained in the Appendix.

2.2 Evidence on Service Outsourcing

Firms have become more specialized over time. As a consequence, more and more processes and components have been handed over to external specialists, contributing to the growth of outsourcing.[12] Although this is a sensible statement, the evidence is quite scattered. Using the Compustat Industry Segment database, Fan and Lang (2000) report some indirect evidence on the increase of specialization; in fact, between 1979 and 1997 the number of publicly traded

[11] In fact Garicano and Rossi-Hansberg (2012b), using a very similar setting, talk more generally about "referral markets".

[12] The definition of outsourcing is standard; in Helpman's (2006) words: "outsourcing means the acquisition of an intermediate input or service from an unaffiliated supplier". This paper will not deal with the choice of the location in which outsourcing is carried out; that is, I will not distinguish between domestic and international outsourcing. The main reason is data limitation but at the same time this paper focuses on service outsourcing for which international outsourcing still plays a relatively little role. For instance Yuskavage et al. (2006) point out that, although the importance of imported services has risen in recent years, their magnitude is still very low, accounting for just 2.7% of total PBS in the U.S. in 2004. Similarly, Amiti and Wei (2009) find that the same share is 2.2% in 2000 and is even lower for other types of services. What this paper will try to shed light on is why firms that are more engaged in trade will have higher shares of domestic service outsourcing.

non-finance firms that operate in a single segment have steadily increased over time. Unfortunately, as pointed out by Antràs and Helpman (2004), there is no systematic analysis on the trend of domestic outsourcing. However, it is possible to get stronger evidence if outsourcing is narrowed to the contracting out of services, and business services in particular. Over the past few decades, firms have purchased more and more services from external providers; namely accounting, engineering, legal services but also security, maintenance, janitorial services just to cite few. This section provides evidence for the aggregate rise of service outsourcing over time.

2.2.1 Industry Level Data

The main reason why the rise of service outsourcing is widely acknowledged is that many of the services that have been intensively contracted out are classified within Professional and Business Services (PBS), and this sector has experienced a dramatic increase over the past few decades. In France, the share of PBS in total employment was 5.4% in 1970, while the same share was 14.7% in 2007, almost a threefold increase. To give a sense to the magnitude of these numbers, consider that the employment share of the total service sector (including the government) has experienced an increase of 20 percentage point, rising from 65.3% in 1970 to 85.2% in 2007, as displayed in Figure 2.3 (left-hand side axis). This is a well-known fact in the structural transformation literature but what has not been sufficiently appreciated is that PBS account for a very large share of this increase. Figure 2.3 also shows the total growth of the service sector and its components (right-hand side axis). PBS have increased their share in total employment by 9.4 percentage points, accounting for 47.2% of the total growth of the entire service sector, the biggest contribution among all industries. Adding Finance, Real Estate and Health Care, these four industries account for almost the entire increase of the service sector in total employment.[13]

The striking rise of PBS would not be sufficient per se to justify an increase in outsourcing. In fact this rise could be driven by final demand. But the PBS sector is quite unusual in this regard: in 2005 roughly 94% of its output was used by firms, either as intermediate inputs or in the form of investment, highlighting the primary role played by firms in this rise. One of the implications of these characteristics is that the remarkable growth in the share of PBS is reflected in a parallel change of the input-output structure of the economy; a fact that has been overlooked in the literature despite the widespread use of input-output data. One way to show this change is looking at the horizontal sum of the coefficients in the total requirements table, usually referred to as forward linkage. This is a measure of the interconnection of a sector to all other sectors through the supply of intermediate inputs. Figure 2.4 shows, for some selected industries, the evolution of the forward linkage divided by the total number of sectors. The figure confirms that PBS have experienced a sharp increase in their forward linkage, overcoming sectors with a traditionally high forward linkage like transportation. PBS have in fact become

[13] The pattern is by no means specific to the French economy, a very similar picture holds true for the U.S., the U.K., and many other developed countries. In Chapter 1, I show that this increase in vertical specialization has a sizable impact on the reallocation of labor across sectors in the U.S., with business services outsourcing alone accounting for 14% of the total increase of the service sector. In ongoing research, I show that the same pattern holds true for most OECD countries and I discuss the impact of these changes on aggregate productivity, see Berlingieri (2013) and Appendix 1.D.

the sector with the highest influence on the rest of the economy, considerably higher than the influence of the average or median sector. In light of the insights provided by Acemoglu et al. (2012), the sharp rise of the PBS forward linkage implies that this sector has greatly increased its influence on the rest of the economy. This fact highlights once more the importance of PBS and why it is key to investigate the reasons that led firms to outsource a higher share of these inputs.

The identification of outsourcing with PBS is quite common in the literature.[14] Yet this assumption could be a source of concern given that industry level data do not clearly distinguish the boundary of the firm, and some of the increase could come from transactions between establishments of the same firm. In Chapter 1, I show that the amount of purchased business services that are reported in the input-output tables are a reliable measure of outsourcing. In fact I control for headquarter establishments and note that the share of internal production remains remarkably constant over time. In any case, in this paper, I overcome these issues looking directly at micro-data, where I can observe business services directly purchased by firms. The downside is that I observe a more limited range of services over the period, with more detailed information available in 2005 only.

2.2.2 Anecdotal Evidence and the Determinants of Service Outsourcing

The evidence on the rise of service outsourcing outlined in the previous section brings about an immediate question: why have firms increasingly contracted out services? And in particular what are the determinants of PBS outsourcing? In order to answer these questions it is insightful to look at some anecdotal evidence first.

An interest case is the experience of Ducati. This firm has been growing very rapidly in recent years, more than tripling the number of bikes produced and expanding to many new markets. Ducati today exports to more than 61 countries. Yet, this success has come with growing pains; among them, inefficiencies in coordinating the production of user manuals and technical documentation, which had to be translated in all the languages of the destination markets. Ducati has therefore decided to contract out its document management to Xerox, which claims to have reduced printing and publishing costs by roughly 20%, together with paper consumption and energy costs. Lowering the costs was certainly a key objective but what managers at Ducati had in mind when they took this decision is probably better represented by the advertisement campaign built on this case. A motorcyclist on a Ducati bike is awkwardly trying to deliver documents inside an office, and the ad goes: "We focus on translating and delivering Ducati's global publications...Which leaves Ducati free to focus on building amazing bike". Another possibly more important objective was therefore to avoid the costs of coordinating all of these peripherical tasks that were stealing the very precious time of managers. Also because the managers could not even monitor the quality of the produced services because they could not certainly learn more than 60 different languages.

[14] Among others, see Abraham and Taylor (1996), Fixler and Siegel (1999), ten Raa and Wolff (2001) and Abramovsky and Griffith (2006).

It is also interesting to note that most service providers like Accenture, KPMG, IBM, McKinsey, Xerox etc... are large multinationals with offices in many countries in the world. This offers a simple explanation for why most of these services are outsourced domestically rather than internationally. Essentially it is likely that these services are "traded" within the borders of these large multinationals. If, for instance, a French firm decides to enter the U.S. market, it is probably going to acquire the service inputs needed like marketing or accounting from the French subsidiary of firms like Accenture or KPMG.

2.3 Existing Literature

2.3.1 Service Outsourcing Literature

Abraham and Taylor (1996) is one of the the very few papers that investigate the determinants of service outsourcing. The authors posit that three main factors may affect the firm's decision to contract out; namely: wage cost savings, the volatility of output demand, and the external provider's specialized skills. The latter consideration refers to the need to access the knowledge and technology provided by the external provider; this comes from the fact that it might not be optimal for a firm to invest in these competencies while an external provider can enjoy economies of scale and amortize the sunk costs of these investments across several clients. Although focused on parts and component production rather than service outsourcing, Bartel et al. (2009) expand this explanation and provide a model in which the probability of outsourcing production is positively related to the firm's expectation of technological change. Investing in a new technology implies some fixed costs; the faster technological change, the shorter the life-span of a new technology, and the less time firms have to amortize their sunk costs. Therefore firms outsource in order to avoid the fixed costs and, at the same time, to access the latest technology possessed by the external providers, which can enjoy economies of scale and spread the fixed costs over a larger demand.

Despite being certainly important, none of these mechanisms can clearly explain why firms that export to more countries, which are usually large firms, outsource a higher share of their costs. Moreover some of the determinants outlined in the previous section have been overlooked, or at least not stressed as the business literature on the other hand does. In particular, the case of Ducati highlights the importance of core competencies and of the managerial challenges that are intrinsically connected with a firm's growth, which often leads to inefficiencies due to complex coordination. Some of these ideas can be found in the Resource-based view of the firm (Penrose, 1959; Wernerfelt, 1984; Prahalad and Hamel, 1990). For instance Quinn and Hilmer (1994) stress that core competencies are skill and knowledge sets and that are usually limited in number: "As work becomes more complex...managers find they cannot be best in every activity in the value chain...they are unable to match the performance of their more focused competitors or suppliers. Each skill set requires intensity and management dedication that cannot tolerate dilution." In linking core competencies to the firm's strategic decision of outsourcing, they also emphasize the role of internal transaction costs and the managerial challenges of producing in-house. These internal transaction costs can be very high and they conclude that: "One of the great gains of outsourcing is the decrease in executive time for managing peripheral activities - freeing top management to focus more on the core of its business."

2.3.2 The Boundaries of the Firm

Most of the literature on the theory of the firm draws the boundary of the firm on the basis of which party owns the asset. But asset ownership is less important in the case of services. For instance, service outsourcing is not very much related to capital intensity, in fact the strong correlation unveiled by Antràs (2003) for good inputs does not hold for services, as shown in Figure 2.5. This is quite intuitive given that services are not capital intensive and there is no reason why the final-good producer should contribute with capital. If anything the production of services is human and knowledge intensive, and the contribution should be in terms of knowledge. But in reality it is quite often the opposite, it is the service provider who has the knowledge on that particular service and a company outsources the service precisely to access that knowledge.

This view is shared by, among others, Rajan and Zingales (2001) who claim that: "as physical assets become less important and give way to human capital, the boundaries of the corporation defined in terms of the ownership of physical assets are becoming less meaningful".[15] The service sector is precisely where "the distinction between ownership and control is important." And services impose a much tighter relationship in the case of integration, which is essentially an employment relationship. This paper therefore embraces a vision of the firm where the residual rights are mainly in terms of control over the decisions to be taken by the firm, and not over the assets; that is, the authority that gives one of the two parties the right to decide the course of action.

2.4 The Model

The model adopts a moral hazard and TCE view of the firm, which stresses the importance of ex-post inefficiencies and specificity even without specific investments. Ex-post adaptation will be at centre stage and the residual rights of control are interpreted as the decision rights to choose the best action in the interest of the organization as a whole. At the same time, the contribution of this paper is to bring back to centre stage the cost of integration as originally stressed by Coase: inside the firm it is the manager who directs and co-ordinates production but there are diminishing returns to management, a given set of activities can hinder the performance of others. Gibbons (2005) stresses the need to explore the complexity of coordination, and the limits that bounded rationality consequently places on firm size and scope. This paper contributes in that direction by modeling integration costs in terms of coordination, following Crémer et al. (2007). Inside the firm the language becomes the coordinator device, highlighting the importance of knowledge and bounded rationality. In fact, the design of a common code captures not only the

[15] In searching a definition for what a firm is, Holmström (1999) adds: "and yet the boundary question is in my view about the distribution of activities: what do firms do rather than what do they own?"

mere communication costs of passing a message but also the larger cognitive costs of interpreting and understanding that message.

2.4.1 Buyer and Suppliers

As common in the trade literature (e.g.: Eaton et al., 2011; Helpman et al., 2008), I assume that some country-specific inputs are needed to export to a given country. It is essentially an Armington assumption on the nature of fixed export costs, for which I will find strong empirical support in the data. In particular, a firm that exports to N countries must source N inputs, one for each destination country.[16] The firm has a simple two-layer hierarchical structure that is fixed: a manager and a certain number of employees. The manager cares about the firm's overall profits and has to decide how to source the inputs. Each input is produced by an agent, and the manager has to make the choice between producing the input in-house by employing the agent directly, or sourcing it from the agent as an external supplier. Moreover there is a trade-off in the way each input is produced. If an agent is focused on producing a certain input i, he will take a very specific action to minimize the cost of producing that particular input, but in this way the input might not fit the overall firm's production process and a coordination cost has to be paid to adapt it. An example could be the production of an accounting software; it can be designed either in a very specific way, with the only objective of recording the transactions of a single product, or in a more flexible way such that it can accommodate the bookkeeping for other products and be linked to the enterprise resource planning system of the firm. In the spirit of Dessein and Santos (2006), I assume a quadratic coordination cost that the manager has to incur if the input is very far from the firm's overall production process.[17]

The production costs that the firm has to incur in order to export to a measure N of countries are given by:

$$C = \int_0^N P_i di + \delta \int_0^N (a_i - \hat{\theta}^m)^2 di + M(N, t, K) \tag{2.1}$$

where P_i is the price paid by the firm for input i, and a_i is the action taken by agent i (employee or external supplier) in order to produce input i. The coordination cost $(a_i - \hat{\theta}^m)^2$ depends on the distance between the action a_i and $\hat{\theta^m}$, the action that would best fit the overall firm's production process (coordinating action); and δ is a parameter that captures the importance of adapting each input to the overall firm's need, hence $\delta(a_i - \hat{\theta}^m)^2$ is the total coordination and adaptation cost for input i. In the baseline version of the model, I take $\hat{\theta^m}$ as a constant, a parameter that characterizes the firm. In Appendix 2.A.2, I propose an extension to the model where the firm is characterized by the actual average action across all inputs ($\bar{a}$). This approach better captures the need of coordinating around the average action that characterizes the firm, introduces extra interesting interdependences across the inputs, and is closer to the TCE idea of ex-post adaptation, given that the mean will depend on the actual realizations of the input

[16] To use calculus and keep the notation simple, I actually write the model in continuum and assume a measure N of inputs. All the qualitative results hold in the possibly more realistic discrete version of the model.

[17] Note, however, that I use a different terminology compared to Dessein and Santos (2006). They use adaptation to mean adapting to the specific local conditions of each task, instead I use adaptation in the classical TCE sense: ex-post coordinated adaptation under hierarchy (e.g. Tadelis and Williamsonn, 2012).

conditions (defined below). The picture that emerges from the extended model is very rich, the firm is characterized by the actual inputs that it needs, and the decision of adding another input (exporting to another country) potentially affects the way in which all other inputs are produced.

Finally $M(N, t, K)$ is the total communication and monitoring costs that the manager has to pay when she decides to employ the agents directly to produce the inputs in-house. These costs are a function of the total number of inputs needed N, the number of employees t, and K, the cognitive ability of the manager.[18] Since each employee produces a single input, t captures both the number of employees and the number of inputs internally produced. The manager has to pay these costs to communicate with her employees and to monitor their actions, which is going to be key in order to be able to steer production inside the firm. Unfortunately the manager is boundedly rational, in the sense that her ability to communicate is going to limited by the maximum number of words she can learn, as in Crémer et al. (2007). Since the firm is identified with the manager who runs it, K is also a source of heterogeneity across firms and in a general equilibrium setting it could be interpreted as the productivity draw of the firm, as in Melitz (2003).

For each input i, there is a market with price taker suppliers and a large pool of entrants. The agents in the market i maximize:

$$\pi_i^s = P_i - (a_i - \theta_i)^2 \geq w_0 \tag{2.2}$$

where a_i is the action that they take to produce input i, and θ_i is the *input condition*, the best way to produce input i. This is essentially the simplest and cheapest way to produce input i separately, without taking in consideration the externalities on coordination costs and, potentially, on the other inputs of the firm. θ_i is a random variable with mean $\hat{\theta}_i$ and variance σ^2. Each input i is therefore characterized by a known distribution with different mean but same variance for all inputs, and the realizations of the input conditions are independent across inputs. Moreover θ_i is private information to agent i, and we shall see how this information will be communicated or not, depending on whether the agent is an employee or an independent supplier. The action a_i is also in general non-contractible (e.g. effort) in the market, while the monitoring activity of the manager will essentially allow her to control the actual action implemented.[19] Finally w_0 is the participation constraint (e.g. the wage that each producer can earn as worker in a non-modeled outside sector), and each market i is ex-ante competitive, so that $\pi_i^s = w_0$ in equilibrium.

2.4.2 Firm Boundaries, Contracts and Timing

The manager can invest in a communication technology that allows her to understand the input conditions and monitor the actions. In this context, the definition of the boundaries of the

[18] More precisely it is a measure N and t of inputs and employees, respectively. Abusing of terminology I use number and measure interchangeably.

[19] In an extension of the model, I also allow for contractible actions in the market (a court can enforce the action) and show that all the qualitative results of the model still go through.

firm are based on the decision to modify the communication technology in order to monitor the actions for input i or not. The definition of integration is not based on the ownership of the asset as in PRT, but it is closer to an employment decision. If the manager decides to produce input i in-house, she will employ the agent in charge of it and will design the communication technology in order to understand the information regarding the input condition θ_i, and to monitor the action of the employee ex-post. This approach fits well the present study given that the inputs are mainly business services, for which physical assets play a relatively small role and the decision of the firm is really whether to employ specialists in-house or not (e.g. having an internal accountant or purchase the accounting services from KPMG).

At time zero, the manager offers a contract, which, in general, is characterized by the tuple $\{P_i, a_i\}$. To avoid confusion I index the set of inputs produced internally (T) by i, and the outsourced ones by j. In the case of outsourcing, the manager cannot contract on a_j because she has not invested in the monitoring technology and the action is also not enforceable in court (this assumption will be relaxed in section 2.A.1). In maximizing his profits according to equation (2.2), the supplier will therefore set $a_j^o = \theta_j$ once the input condition will be realized. The manager has no way to avoid this action because she has not invested in the monitoring technology, and even if she could (e.g. in case a court could enforce a particular action), she would not be able to improve much in terms of coordination because she would not know the actual realization of the input condition (no communication). The best that the manager can do in this situation is to simply offer $P_j^o = w_0$. Therefore, in the case of outsourcing, the contract is characterized by a fixed price: the market gives high-powered incentives to the independent supplier.

On the other hand, in the case of integration (employment), the manager can contract on a_i thanks to monitoring. She will tell the employee to implement a certain action a_i^v and will pay him: $P_i^v = w_0 + (a_i^v - \theta_i)^2$. Therefore employment is characterized by what Bajari and Tadelis (2001) refer to as a C+ contract, a contract that pays a fixed wage plus any cost the agent might incur in producing the input. This is the closest situation to an actual employment contract: the manager has the power and authority to tell the employee what to do but she compensates him of any cost, providing soft-power incentives.

As it will be clearer in the next sub-section, the make or buy decision is driven by the trade-off between the benefits of a better ex post coordination in-house and the costs of investing in the communication technology. In fact by employing the agent directly the manager can learn the actual realization of the input condition and hence achieve a better coordination ex post by internalizing the (negative) externality that the agent's action has on the rest of the organization. At the same time the manager has to compensate the employee so that he will be willing to perform an action that is not strictly optimal for the specific input on its own. Outsourcing, on the other hand, reaches precisely the opposite result: the market offers high power incentives through a fixed price, which will give the incentive to the external supplier to take the cost-minimizing action for that particular input. Outsourcing allows the organization to source that particular input at the ex-ante minimum price. This is of course the best thing to do when the importance of adaptation (δ) is low or when setting up a common communication

Figure 2.1: The timing

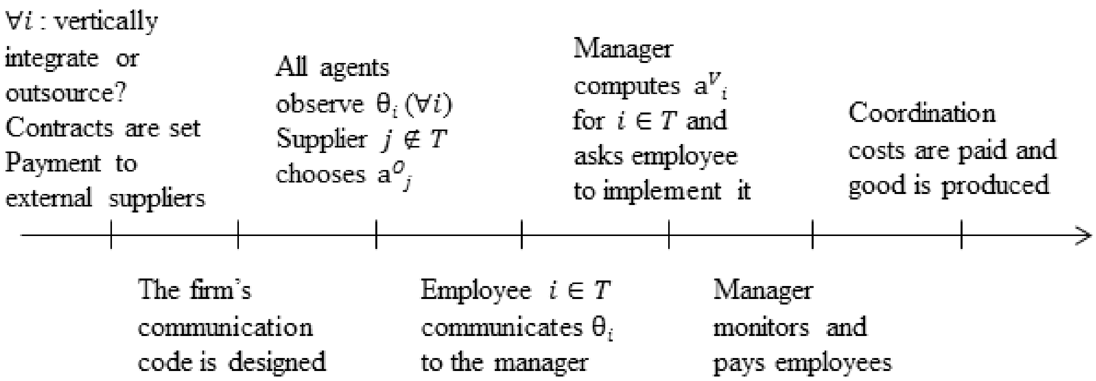

code is very costly, because the firm does not have to design it.[20] The drawback is of course very high coordination costs in case adaptation is needed.

Figure 2.1 clarifies the timing and the details of the game. At the beginning of the period the manager decides for each input whether to source it from an external supplier or to produce it internally by employing the agent. The contracts are set and the external suppliers are immediately paid given that the price is fixed and does not depend on the action taken by the agent ($P_j^o = w_0$). The employee instead will be paid his wage and any cost he has incurred in producing the input ($P_i^v = w_0 + (a_i^v - \theta_i)^2$) after the input condition is realized and he has implemented the agreed action. Once the set of inputs produced internally is decided, the manager designs the communication code that serves two purposes: understanding the messages of the employees when they communicate their input conditions, and monitoring their actions ex-post to check they have performed what they were told to do.

After the code is designed, all agents (both external suppliers and employees) observe their input conditions. At that point the external suppliers take their optimal action, which, given their payoff in (2.2) and the fact that they have received a fixed price, is clearly going to be: $a_j^{o*} = \theta_j$. On the other hand, the employees communicate their input conditions to the manager, who can understand them since the communication code has been designed precisely to interpret the messages for that specific set of inputs (I will describe the communication technology in more detail in Subsection 2.4.4). Note that the employees have the incentive to truthfully communicate their input conditions because they will be compensated for any cost they will incur, and because the actions they will be assigned will depend on the input conditions they have communicated. If an employee deviates and communicates a different input condition, his assigned action will reflect that and not the actual input condition. The manager will monitor the action implemented by the employee so he will have to do precisely what requested.[21] If the actual input condition is different from what the employee has communicated, the total price

[20]Due to the lack data I do not distinguish between domestic and internal outsourcing, but comparing P_j^o and P_i^v it is immediately clear that international outsourcing to a low-wage country (low w_0) is attractive whenever the importance of adaptation is low.

[21]Otherwise he gets punished. For instance he could be fired and receive a zero wage.

(wage and compensation for the costs) will not compensate the actual costs incurred by the employee and he will be worse off.

On the other hand communication is not possible under outsourcing. In this setting the reason is due to the fact that the manager has not designed the communication code to understand the messages related to the input conditions of the outsourced inputs. Any message from the external suppliers is therefore pure white noise for the manager.[22] This is clearly an extreme case but it reflects the fact that internal communication within the firm is usually more effective compared to communication with external suppliers, because the incentives are more aligned. For instance in a full strategic communication setting, Alonso et al. (2008) show that even allowing for some degree of horizontal communication, vertical communication is always more effective. Moreover the present setting resembles the hard but costly communication proposed by Dewatripont and Tirole (2005).

Once the manager has learned the input conditions, she is in the position to compute the optimal actions for all inputs. In doing so, she minimizes the costs of producing each input as well as the coordination costs in case of adaptation. Essentially the manager is capable of achieving coordination at a lower cost because she can internalize the negative externality that each input imposes on the rest of the organization. The manager then tells each employee what to do, monitors that they implement precisely what requested, and pays them. Finally coordination and adaptation costs are paid and the good is delivered. The optimal internal actions are obtained in the next subsection.

2.4.3 Optimal actions

The problem is solved by backward induction. Once input conditions are revealed, the manager chooses the action a_i^v for each input internally produced ($\forall i \in T$) in order to minimize total costs. Assuming that a measure t of inputs is internally produced by an equal measure of employees, the problem of the manager is the following one:

$$\min_{\{a_i^v\}} Nw_0 + \int_0^t (a_i^v - \theta_i)^2 di + \delta \int_0^t (a_i^v - \hat{\theta}^m)^2 di + E\left[\delta \int_t^N (a_j^o - \hat{\theta}^m)^2 dj\right] \tag{2.3}$$

It is easy to show that the optimal action is a weighted average of input condition and the coordinating action:

$$a_i^{v*}(\theta_i, \hat{\theta}^m) = \frac{1}{1+\delta}\theta_i + \frac{\delta}{1+\delta}\hat{\theta}^m \tag{2.4}$$

The manager internalizes the externality that each input imposes on coordination costs when adaptation is needed. Therefore the optimal action will lie somewhere in between the action that minimizes the production costs of the specific input and the coordinating action that best fits the overall firm's production process and minimizes coordination costs in case of adaptation. In this simple baseline model, the optimal action for input i does not depend on the actions for other inputs. There will only be some limited interaction across inputs from the communication

[22] Even if the manager could infer something, we would not be able to influence the action taken by the external supplier.

Figure 2.2: A Simple Representation

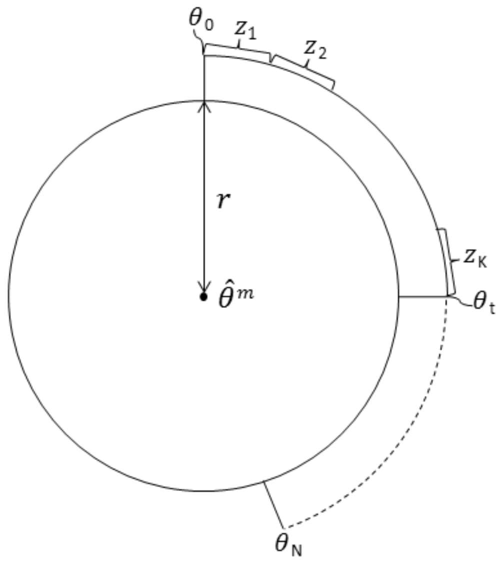

costs; on the other hand, in the extension of Appendix 2.A.2, the optimal action for each input will depend on all other internal actions, providing a full interaction across inputs.

Since the external suppliers take an action $a_j^{o*} = \theta_j$ (minimizing costs for that particular input), and exploiting the fact that all input conditions are independently drawn from distributions with the same variance, it is fairly easy to show that the expected costs at time zero are:

$$E[C] = Nw_0 + \left[\frac{\delta}{1+\delta}t + \delta(N-t)\right]\sigma^2 + \\ + \frac{\delta}{1+\delta}\int_0^t (\hat{\theta}_i - \hat{\theta}^m)^2 di + \delta\int_t^N (\hat{\theta}_j - \hat{\theta}^m)^2 dj + M(t,N,K) \quad (2.5)$$

A simple way of representing the problem is to assume that all input conditions lie on a circle and that the coordination action $\hat{\theta}^m$ is at the center (Figure 2.2). Each input condition is a point on the circle that has a length of measure $\bar{N}$, which can be interpreted as the maximum measure of countries the firm can export to (i.e. number of countries in the world). In this way the distance between the coordination action and the means of the input conditions is the same for all inputs and is simply pinned down by the length of the circle ($r = \bar{N}/2\pi$). Under this simple representation, the expected costs at time zero can be further simplified as follows:

$$E[C] = Nw_0 + \left[\frac{\delta}{1+\delta}t + \delta(N-t)\right](\sigma^2 + r^2) + M(t,N,K) \quad (2.6)$$

This expression shows very clearly the trade-off between outsourcing and integration. The second term is decreasing in t and captures the benefits of integration: by coordinating the actions in-house the manager is able to steer production, internalize the externalities, and achieve a lower coordination costs. On the other hand, by producing more inputs in-house, the communication costs (third term) intuitively rise. In this simple baseline model, the returns to in-sourcing do not depend on t, and the communication/monitoring costs do not depend on the importance of adaptation δ and the variance of the input conditions σ^2. Therefore even without specifying much of the nature of the communication costs, I can state the first Proposition of the paper.

PROPOSITION 1: If the communication costs increase in t ($\frac{\partial M(t,N,K)}{\partial t} > 0$), the expected profits $E[\pi] = -E[C]$ are supermodular in t, σ^2, δ. Hence, the optimal number (measure) of inputs produced in-house t^* increases with the importance of adaptation (δ) and the volatility of the input conditions (σ^2): $\frac{\partial t^*}{\partial \delta} > 0$ and $\frac{\partial t^*}{\partial \sigma^2} > 0$

COROLLARY 1: the complementarity with the variance of input conditions disappears when the importance of adaptation goes to zero.

$$\lim_{\delta \to 0} \frac{\partial^2 E[\pi]}{\partial t \partial \sigma^2} = \lim_{\delta \to 0} \frac{\delta^2}{1+\delta} = 0 \tag{2.7}$$

The proof is immediate and follows standard results on supermodularity.[23] The intuition is also quite straightforward. If adaptation is not very important, there is clearly no reason to produce the inputs in-house because it is true that the manager can achieve better coordination but this is needed only in case of adaptation and is costly due to communication costs.[24] If the conditions are very volatile, a situation in which the firm does not really know what it gets or inputs are not very homogeneous, the firm will find it optimal to in-source more in order to reduce the risk of having an input that will be very costly to coordinate because very far from $\hat{\theta}^m$. Intuitively Corollary 1 further specifies that this effect clearly disappears when adaptation is not very important. These results are quite general because rely on a very simple and easily satisfied assumption, namely that the communication costs increase in the number of inputs internally produced.

2.4.4 Communication and Monitoring Costs

The manager has to communicate with all employees to learn their input conditions and then monitoring their actions. Therefore the total monitoring costs are given by: $M(t,N,K) = \int D(t,N,K)dt$, where $D(t,N,K)$ is the total diagnosis cost that the manager has to pay to understand the message and monitor each employee. Following Crémer et al. (2007), the manager is boundedly rational and can learn K words at most. Each word allows the manager to identify and recognize a certain set of input conditions, which is referred to as the breath of the word. The more general the word is (wider breadth), the higher is the diagnosis cost that the manager has to incur in order to understand the content of the message. I simply assume that the diagnosis cost for each word is linear in its breadth. Therefore the total expected diagnosis cost is given by: $D(t,N,K) = \sum_{k=1}^{K} p(z_k) z_k$, where z_k is the breadth of word k and $p(z_k)$ is its probability.

[23] See for instance Milgrom and Roberts (1995).

[24] An interesting extension to the model would be to have an heterogeneous adaptation need across inputs, say δ_i, which could be interpreted as the probability of adaptation for that particular input. In this case it is intuitive that the first inputs to be outsourced would be the ones with very low adaptation probabilities, that is the inputs that are fairly homogeneous or that are very far from the core competency of the firm, so that the firm has no need to adapt them to its own production process.

The words are of course not overlapping, that is each of them refers to different sets of events to minimize the costs.

In the simple setting of the present model, all the inputs are needed in equal proportions (the production technology is Leontief), hence the manager will face the same probability of communicating with each employee. Essentially the manager has to communicate with all employees and therefore she will encounter all the input conditions with equal probability. In the simple representation of Figure 2.2, each input condition lies on a point of the circle, and the manager will therefore face an overall uniform distribution of events. Each input condition is in fact equally likely and the overall distribution of input conditions that the manager encounters is an uniform on the interval $[0, N]$, where N is again the measure of total input needed or, in the empirical interpretation, the total number of countries the firm is exporting to.

The manager will design an optimal code to minimize the total expected diagnosis cost given the inputs that are produced internally. She will solve the following problem:

$$\min_{\{z_k\}_{k=1}^{K}} \sum_{k=1}^{K} p(z_k) z_k \quad s.t. \quad \sum_{k=1}^{K} z_k = t \tag{2.8}$$

where $p(z_k) = \frac{z_k}{N}$ due to the fact that the underlying events are uniformly distributed. In this setting the solution to the problem is very simple and all the words have the exact same breath: $z_k = z_h = \frac{t}{K}$, $\forall l, h$. Therefore the total expected diagnosis cost for each employee is given by: $D(t, N, K) = \frac{t^2}{KN}$. And the total communication costs are defined as follows:

$$M(t, N, K) = \frac{t^3}{3KN} \tag{2.9}$$

The costs are intuitively decreasing in the cognitive ability of the manager K and depend on the set of inputs that are internally produced. Adding more inputs raises the costs for all other inputs already internally produced because the manager has to change the code, and make all words more imprecise in order to accommodate the new set of input conditions. Finally, since the optimal actions refer to the same underlying set of input conditions, the communication technology also allows the manager to monitor the actions of the employees (for simplicity I assume that the cost is paid only once for both activities).

2.4.5 The Optimal Outsourcing Share and the Effect of Globalization

Assuming the previous form of communication/monitoring costs, it is easy to solve the problem that the manager faces at time zero and find the optimal measure of inputs internally produced. A simple minimization of the expected costs in (2.6) with the communication costs as in (2.9), gives the optimal share of inputs internally produced:

$$\frac{t^*}{N} = \delta \sqrt{\frac{K\psi^2}{(1+\delta)N}} \qquad \textit{where: } \psi^2 = \sigma^2 + r^2 \tag{2.10}$$

I can therefore state the two main propositions of the paper, that will be tested in the empirical section.

Proposition 2: If the number (measure) of destination countries increases, $N \uparrow$, the optimal share of outsourced inputs increases:

$$\frac{\partial}{\partial N}\left(1-\frac{t^*}{N}\right)=\frac{1}{2}\frac{t^*}{N}\frac{1}{N}>0 \tag{2.11}$$

Proposition 3: the share of outsourced inputs is concave in N:

$$\frac{\partial^2}{\partial N^2}\left(1-\frac{t^*}{N}\right)=-\frac{3}{4}\frac{t^*}{N}\frac{1}{N^2}<0 \tag{2.12}$$

Proposition 2 gives the main effect of interest that will be extensively investigated in the empirical part of the paper. When the number of destination countries increases, so does the number of inputs needed to reach those destinations, making the coordination of these inputs more and more complex. The reason is that the manager has to design a code that needs to accommodate a larger set of different events, all arising with a very small probability. This makes communication inside the firm very costly. The coordination benefits are still present and the absolute number of inputs internally produced still increases, but their share in the total number of inputs decreases. The reason is that the manager, facing too high communication costs, finds it optimal to outsource to get the benefits of a low ex-ante price. Moreover Proposition 3 shows that the relationship between the optimal share of outsourced inputs and the number of destination countries is non-linear, and concave in particular.

Another interesting and intuitive result is captured by the following proposition:

Proposition 4: A manager with higher cognitive ability (K - measure of skill) in-sources a higher share of inputs: $\frac{\partial(t^*/N)}{\partial K}>0$

It is also interesting to note that the total expected costs of the firm are decreasing in the manager's skill (K) and increasing in the importance of adaptation (δ):

$$E[C]=Nw_0+\delta N\psi^2-\frac{2}{3}\frac{\delta^3\psi^3}{(1+\delta)}\sqrt{\frac{NK}{(1+\delta)}} \tag{2.13}$$

And the increase in total expected costs to export to an extra country is given by:

$$\frac{\partial E[C]}{\partial N}=w_0+\delta\psi^2-\frac{1}{3}\frac{\delta^3\psi^3}{(1+\delta)}\sqrt{\frac{K}{(1+\delta)N}} \tag{2.14}$$

The marginal cost of exporting to an extra country is therefore increasing in N. In a setting with a discrete number of countries, this would correspond to the fixed cost to export to an extra country. The simple setting of the baseline model has therefore the potential to be nested in a general equilibrium trade model (e.g. Melitz, 2003, with productivity draws on the level of managerial ability) to have a full set of trade implications. This extension is left for future research.

2.4.6 A General Condition on the Communication/Monitoring Cost Function

Another interesting extension of the paper is to study the general set of communication and monitoring cost functions that are consistent with the empirical results. The optimal number of inputs produced in-house is pinned down by:

$$\frac{\delta^2}{1+\delta}(\sigma^2 + r^2) - M_t(t, N, K) = 0 \tag{2.15}$$

where $M_t(t, N, K)$ is the marginal communication cost with respect to an increase in the number of inputs internally produced. The main finding in the empirical results corresponds to the following condition:

$$\frac{\partial t^*(N)/N}{\partial N} < 0 \Longleftrightarrow \varepsilon_{t^*,N} < 1 \tag{2.16}$$

where $\varepsilon_{t^*,N}$ is the elasticity of t^* with respect to N. The problem is separable and the previous condition boils down to a constraint on form of the monitoring function:

$$\varepsilon_{t^*,N} < 1 \Longleftrightarrow -\frac{\varepsilon_{M_t,N}}{\varepsilon_{M_t,t}} < 1 \tag{2.17}$$

where $\varepsilon_{M_t,N}$ and $\varepsilon_{M_t,t}$ are the elasticities of the marginal communication cost with respect to the number of countries and the number of inputs internally produced, respectively.[25]

It is interesting to note that the result is certainly more general than the setting of the baseline model, which gives a simple expression for the effect of interest but is based on the specific case of an uniform distribution for the overall events that the manager faces. Drawing on the intuition provided by Crémer et al. (2007), we know that an uniform distribution of events is actually the worst case, with the highest level of communication costs. The reason is that all events are equally likely and the manager cannot design a code targeted to a certain set of more frequent events. In a more general setting, the optimal code features words of different breadths: in order to save on the diagnosis costs, very precise words are used to refer to very frequent events, while very broad and costly words are used for rare events. Still, if fixed export costs are country specific, adding another country implies adding different events or events that were before very unlikely. Hence the overall distribution will tend to get closer to a uniform distribution when the firm will start exporting to more and more countries.

2.5 Econometric Evidence from France

2.5.1 Data

The model is tested using firm level data from France for the period 1996-2007. I rely on four main data sources. First, the Enquête annuelle d'Entreprise (EAE) that collects balance sheet data on all French firms with more than 20 employees and a sample of smaller firms. Second, the

[25] The problem is not separable and this condition does not hold only in the discrete version of the extension presented in Appendix 2.A.2.

Déclaration annuelle de données sociales (DADS) that collects employment data on all firms with paid employees; the data used are aggregated at the establishment level. Third, transaction level import-export data come from the French Customs; these data have been used among others by Eaton et al. (2004). Finally, service outsourcing data contained in the EAE are integrated with the Enquête Recours aux Services par l'Industrie (ERSI), a survey of firms with more than 20 employees and the census of firms with more than 250 employees that collects detailed information about service outsourcing policies for the year 2005. The analysis will mainly focus on manufacturing firms (NACE Rev1.1 D category).

Table 2.1: Summary Statistics by Export Status - 2005

	Nonexporters			Exporters		
	Mean	Median	N	Mean	Median	N
Employment	44.6	30	5,220	158.4	48	16,453
Turnover	7,107	3,331	5,076	47,257	8,577	16,360
Total Exports	0	0	5,307	13,575	793	16,497
Num. Countries	0	0	5,307	13.5	7	16,497
Num. Products	0	0	5,307	14.7	7	16,497
Num. Languages	0	0	5,307	9.44	7	16,497
K/L Ratio	52.8	23.5	5,057	99.1	43.4	16,336
S/U Ratio	0.65	0.26	4,984	1.12	0.42	15,961
Professionals Sh.	0.074	0.045	5,049	0.13	0.086	16,171
HQ Intensity	0.035	0	5,031	0.069	0	16,306
PBS Out. Sh. in Costs - 1	0.034	0.0045	4,800	0.047	0.013	15,951
PBS Out. Sh. in Costs - 1b	0.023	0.0037	4,800	0.034	0.01	15,951
PBS Out. Sh. in Costs - 2	0.034	0.0046	4,800	0.047	0.014	15,951
PBS Out. Sh. in Costs - 3	0.034	0.0049	4,934	0.048	0.015	16,241

Note: Turnover, total exports, and K/L ratio are measured in thousands of €. Full sample.

Table 2.1 reports summary statistics for the main variables in 2005, separately by export status. The EAE mainly contains large firms, so, not surprisingly, the majority of firms in the sample are exporters.[26] As well known from the trade literature, exporters are larger, and more capital and skilled intensive. The average exporter in the sample exports 14.7 products to 13.5 destination countries, which sometimes share the same language, in fact the number of destination languages is 9.5. Moreover exporters outsource more professional and business services. Considering the baseline definition of business services (PBS Out. Sh. in Costs - 1) exporters spend the equivalent of 4.7% of their total costs in business services purchased from the market, compared to 3.8% for non-exporters.[27] In the baseline definition business services outsourcing includes: purchases of studies, expenses related to the purchase of IT services, and advertisement. I also propose alternative measures in which I add non-capital expenditures on software purchases (measure 2), and capital expenditures on software purchases and investment

[26]The firms in the sample account for 87.5% of the total turnover of the French manufacturing sector in 2005 (aggregate data from Eurostat).
[27]Note that exporters also produce more services in-house: HQ Intensity is in fact higher.

in R&D (measure 3).[28] More precise variable definitions and the procedure employed to clean the data are described in Appendix 2.B.

Table 2.2 shows the change over time for the main variables of interest. On average firms have increased their share of outsourced services in total costs by 10%, from a share of 3.86% in 1996 up to 4.25% in 2007. The average firm has increased the number of export destination countries from 7.9 in 1996 to 10 in 2007, equivalent to a 27.5% increase.

Table 2.2: Change in Outsourcing Shares and Destination Countries

	1996	2007	Change
PBS Out. Sh. in Costs - 1	0.0386	0.0425	10.10%
PBS Out. Sh. in Costs - 2	0.0386	0.0426	10.36%
PBS Out. Sh. in Costs - 3	0.0397	0.0432	8.82%
Num. Countries	7.8787	10.0427	27.47%

2.5.2 The Impact of Coordination Complexity on PBS Outsourcing

By averaging across all firms exporting to a certain number of markets in all years, Figure 2.6 shows that the share of purchased business service on sales is positively and significantly related to the number of export destination countries, the main measure of coordination complexity used in the analysis. The simple intuition is that the higher the number of countries a firm is exporting to, the more complex its business environment is going to be. This is very much in line with the most common definition of complexity in systems theory, where complexity arises through connectivity and the inter-relationships of a system's constituent elements. In the present case, the higher the number of connections (destination countries), the higher coordination complexity is going to be, because exporting requires more inputs. Designing a communication code for all these infrequent events is very costly and therefore, according to Proposition 2, the share of outsourced inputs in total costs increases.

The simple correlation for the average firm is confirmed when the full panel of firm-level data is analyzed. I run the following simple reduced form regression:

$$OUT_{it} = \beta_0 + \beta_1 NC_{it} + \delta + \epsilon_{it} \tag{2.18}$$

where OUT_{it} is the share of purchased business services over total cost for firm i at time t, NC_{it} is the number of export destination countries, and δ is a set of fixed effects. Proposition 2 predicts that β_1 should have a positive sign. In the baseline regression, business services are measured as the sum of: purchases of studies, expenses related to the purchase of IT services, and advertisement. Table 2.3 shows the results of the regressions. Column (1) shows that the export status of the firm is positively and significantly related to the share of purchased business services, that is, the trade extensive margin is positively related to service outsourcing. Column (2) shows

[28] The latter measure is probably the less reliable because it is not possible to completely rule out the possibility that part of the R&D investment is actually performed in-house.

that coordination complexity, measured as the number of export destination countries, has a strongly positive and significant effect on the share of purchased business services. Therefore, among exporters, the firms that export to more countries tend to outsource a higher proportion of services. This hitherto unknown systematic pattern is actually fairly intuitive and goes well with the existing literature on international trade. The fixed export costs are often characterized as the specific service inputs needed to export to a particular country; hence exporting to more destination countries implies that more inputs are needed (e.g.: a different advertising campaign for each destination market).[29] Each of these country-specific service inputs is a low probability event from the point of view of the manager of the manufacturing firm; and if a firm exports to more countries, each of these events becomes even less frequent, which translates into a more complex business environment. The model predicts that the share of outsourced inputs in total costs increases because the firm has no incentive to invest in the communication technology to produce these inputs in-house: the presence of these very infrequent inputs makes communication and monitoring very costly.

Table 2.3: Purchased Business Services and Coordination Complexity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Var.	Out. 1	Out. 1	Out. 1	Out. 1	Out. 1	Out. 2	Out. 3	Out. 1
Exporter	0.554***		0.098***					
	(0.020)		(0.017)					
NC		0.185***		0.083***	0.086***	0.087***	0.078***	0.086***
		(0.008)		(0.010)	(0.010)	(0.010)	(0.010)	(0.022)
Observations	235,182	184,556	235,182	184,556	184,556	184,864	186,725	184,556
Num. of firms	39,500	31,212	39,500	31,212	31,212	31,246	31,380	31,212
Year FE	Yes	Yes	Yes	Yes				
Industry FE	Yes	Yes						
Industry-Year FE					Yes	Yes	Yes	Yes
Firm FE			Yes	Yes	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Industry

Note: The dependent variable is the share of purchased services over total costs measured in logs. In columns (1)-(5) and (8) business services are measured as the sum of purchases of studies, expenses related to the purchase of IT services, and advertisement. Column (6) adds non-capital expenditures on software purchases. Columns (7) adds capital expenditures on software purchases and investment in R&D. Data are for period 1996-2007. Clustered standard errors in parentheses; (*, **. ***) indicate 10, 5, and 1 percent significance levels.

The main focus of the analysis is actually the increase in outsourcing of business services over time and an obvious worry is that firms differ across a variety of other dimensions. Hence from columns (3) onwards I add firm fixed effects to focus on the within firm variation and control for unobserved time-invariant firm characteristics. Another worry is that industry specific shocks that occur in some periods in some industries might affect the results, hence from column (5) onwards I control for a full set of interacted industry-by-year fixed effects. By focusing

[29]In motivating the presence of some fixed costs to exporting, Melitz (2003) asserts that a firm must inform foreign buyers about its product, learn about the foreign market, research the foreign regulatory environment etc... These tasks correspond to advertising, market and legal research, and they are all supplied by the professional and business industry. Das et al. (2007) and Morales et al. (2011) put forward very similar arguments. Among others, Eaton et al. (2011) and Helpman et al. (2008) adopt settings that feature country-specific fixed export costs.

on the within variation the magnitude of the results is smaller but becoming an exporter or increasing the number of export destination is still highly related to more business services outsourcing.[30] The effect is also quantitatively important. Considering the coefficient in column 5 and the variation shown in Table 2.2, for the average firm in the sample, the increase in the number of destination countries explains around 20% of the increase in the share of outsourced services over the period. The following columns show that very little changes when I modify the measure of outsourcing by adding non-capital expenditures on software purchases (column 6), and capital expenditures on software purchases and investment in R&D (column 7). Finally column (8) shows that the result is still highly significant even after clustering standard errors at the industry level and performing the full degrees of freedom adjustment due to non-nested panels within clusters (firms that change industry).

The literature on firm boundaries has proposed many other potentially time-varying determinants that could affect outsourcing, other than the proposed proxy for coordination complexity. I therefore modify the basic regression to include other controls, I run the following regression:

$$OUT_{it} = \alpha_i + \beta_1 NC_{it} + \boldsymbol{W'_{it}\beta_2} + \delta_{jt} + \epsilon_{it} \tag{2.19}$$

where OUT_{it} and NC_{it} are defined as before, $\boldsymbol{W_{it}}$ is a vector of controls, α_i are firm fixed effects, and δ_{jt} are the full interaction of industry and year fixed effects. The first interesting question to ask is whether outsourcing of services is a fixed cost component as assumed so far, or it also entails tasks related to variable costs. In the latter case outsourcing of business services would be also affected by the exporting intensive margin. I therefore include total exports in the regression. As column (2) of Table 2.4 shows, the intensive margin is in fact not driving the purchase of business services. Hence the common assumption that business services are a fixed export cost component seems to hold in the data. I will further investigate this issue in Section 2.5.4.

Table 2.4 also shows the baseline results after the inclusion of several controls common in the vertical integration literature: capital intensity, human-capital intensity, a measure of scale economies, value-added over sales, as proposed by Antràs (2003); a measure firm-level contract intensity in the spirit of Nunn (2007); and a measure of headquarter intensity as proposed by Antràs and Helpman (2004, 2008).[31] The effect of coordination complexity remains robust and stable to the inclusion of all controls. Most importantly column (6) includes a measure of internal production of services. The model predicts that when the number of infrequent tasks increases, these will be outsourced. It is therefore important to verify that this increase in the need of service inputs is not driven by an overall re-focus of the firm towards service activity, which would imply a parallel increase in the internal production of services. Controlling for internal production is not easy, even with micro-level data, because it is very hard to observe the tasks internally produced by firms. In Chapter 1, using data for the U.S., I show that this is true at the aggregate manufacturing industry level, after controlling for business services produced

[30] A simple reason for the smaller magnitude could be the classic attenuation bias from measurement error, see Angrist and Pischke (2008a).

[31] Results are also robust to the inclusion of a measure of TFP.

in-house. I adopt here a similar strategy and I control for the share of revenues generated by establishments of the firm that are classified within services, essentially I control for the revenue share of headquarters. In this specification I include establishments producing any type of services, not only business services. As column (6) shows the results are very robust to this measure of internal production.

Table 2.4: Purchased Business Services and Coordination Complexity - Covariates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Exporter	0.098***								
	(0.017)								
NC		0.086***	0.082***	0.083***	0.086***	0.086***	0.080***	0.085***	0.089***
		(0.010)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.013)	(0.013)
Exports			0.003	0.003	0.001	0.002	-0.001	-0.000	0.000
			(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)
Capital Intensity				0.030***	0.028**	0.028**	0.035***	0.024*	0.022
				(0.011)	(0.012)	(0.012)	(0.013)	(0.014)	(0.014)
Skill Intensity					0.047***	0.046***	0.051***	0.051***	0.051***
					(0.010)	(0.010)	(0.010)	(0.011)	(0.012)
HQ Intensity						-0.052	-0.054	-0.047	-0.054
						(0.037)	(0.037)	(0.039)	(0.039)
Scale							0.081***	0.066***	0.065***
							(0.022)	(0.023)	(0.024)
Num imp. products								0.014*	0.019**
								(0.008)	(0.009)
Contract Intensity									0.015**
									(0.006)
Observations	235,182	184,556	184,556	183,487	174,908	174,682	174,682	150,874	144,927
Number of firms	39,500	31,212	31,212	31,073	30,172	30,159	30,159	26,091	25,339
R-sq W	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of purchased services over total costs. All variables are in logs apart from HQ Intensity. Data are for period 1996-2007. Firm-clustered standard errors in parentheses; (*, **. ***) indicate 10, 5, and 1 percent significance levels.

Some of the other controls are also worth discussing. It is interesting to note that the firm scale (total number of employees) and the total number of imported goods are positive and significant, even though not very robust. Both variables could be interpreted as alternative measures of managerial complexity. Capital intensity is positive, although not always significant. This implies that firms that increase their capital stock are more likely to outsource business services. Moreover the positive and significant effect of the contract intensity variable can be interpreted as another support, albeit indirect, to the complexity and core-competencies story. The variable is constructed using the information about firms' imports. The firm-level contract intensity is therefore a weighted average of the contract intensity of all firm imports, where the measure of contract intensity is taken from Rauch (1999), analogously to Nunn (2007), and the weights are the shares of each product in the total firm imports. Under a standard TCE interpretation, as also pointed out by Corcos et al. (2013), a firm in-sources more contract intensive inputs. Given that all of the observed imports are goods, the positive impact on service outsourcing can be rationalized by arguing that a manufacturing firm with more contract intensive inputs focuses on its core-competencies by producing more goods in-house and outsourcing more of the non-core services.

2.5.2.1 Alternative Measures of Complexity

The multi-product literature assumes the presence of product-specific export fixed costs, and again these are arguably mainly made up by service inputs. For instance Bernard et al. (2011) justify the presence of product-specific fixed costs arguing that they capture the research, advertising, and regulation costs to supply each product to a certain destination. Therefore also an increase in the number of products could entail an increase in the number of service inputs needed, and consequently an increase in coordination complexity. Table 2.5 shows the results when also the number of products and the interaction between number of products and countries (demeaned) are added. There is indeed a positive and significant relationship between the number of exported products and the share of service outsourcing. The magnitude is smaller compared to the number of destination countries. This result is in line with the fact that the fixed costs to export to a new destination are higher than those needed to export a new product, and that the market-specific entry costs drop fast with the number of products, as shown by Arkolakis and Muendler (2010). As expected, the interaction between the two variables is positive as well.

Table 2.5: PBS Outsourcing and Number of Products

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
NC	0.086***	0.068***	0.082***	0.080***	0.081***	0.084***	0.083***	0.078***	0.081***	0.084***
	(0.010)	(0.011)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.013)	(0.014)	(0.014)
NP		0.028***	0.038***	0.036***	0.036***	0.036***	0.037***	0.034***	0.032***	0.031***
		(0.008)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.010)	(0.010)
NC#NP			0.029***	0.030***	0.030***	0.030***	0.030***	0.027***	0.027***	0.028***
			(0.005)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Exports				0.003	0.002	0.001	0.001	-0.001	0.000	0.001
				(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)
Capital Intensity					0.031***	0.029**	0.029**	0.035***	0.024*	0.022
					(0.011)	(0.012)	(0.012)	(0.012)	(0.014)	(0.014)
Skill Intensity						0.047***	0.046***	0.050***	0.050***	0.050***
						(0.010)	(0.010)	(0.010)	(0.011)	(0.012)
HQ Intensity							-0.058	-0.059	-0.051	-0.059
							(0.037)	(0.037)	(0.038)	(0.039)
Scale								0.064***	0.051**	0.049**
								(0.022)	(0.023)	(0.024)
Num imp. products									0.008	0.013
									(0.008)	(0.009)
Contract Intensity										0.015**
										(0.006)
Observations	184,556	184,556	184,556	184,556	183,487	174,908	174,682	174,682	150,874	144,927
Number of firms	31,212	31,212	31,212	31,212	31,073	30,172	30,159	30,159	26,091	25,339
R-sq W	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of purchased services over total costs. All variables are in logs and interaction variables are demeaned. See notes in Table 2.4.

The proxy of coordination complexity defined as the number of destination countries is appealing for its simplicity. At the same time it is a rather crude measure, I therefore propose few other possible ways of measuring complexity. Arguably the majority of the fixed export costs could be related to handling transactions in a different language (translating labels, instructions, advertising campaigns, different legal system etc...).[32] Instead of counting the number of countries I therefore count the number of different languages. More generally, as shown by

[32]The gravity literature has shown that sharing a common language is a trade facilitator. A different language can also capture deeper cultural barriers.

Table 2.6: Alternative Measures of Coordination Complexity

	(1)	(2)	(3)	(4)	(5)	(6)
Exports	0.003	0.010	0.004	0.007	0.006	0.004
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
NC	0.135***					
	(0.016)					
Num Languages		0.113***				
		(0.014)				
NC (Gravity)			0.123***			
			(0.016)			
NC (WB Doing Business)				0.100***		
				(0.014)		
NC (WB Doing Business - Trade)					0.100***	
					(0.014)	
NC (Complexity)						0.143***
						(0.016)
Observations	176,492	176,492	176,492	176,492	176,492	176,492
Number of firms	30,438	30,438	30,438	30,438	30,438	30,438
R-sq W	0.01	0.01	0.01	0.01	0.01	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of purchased services over total costs. All columns also include the following regressors: K/L, S/L, HQ Intensity, and Scale. Regressors are standardized. Only exporters are included in the sample. See also notes in Table 2.4.

Morales et al. (2011) in their extended gravity framework, the fixed costs of exporting could be related to the standard gravity variables like common language, continent, and legal system. I therefore give countries a weight of 1 if they share all of the previous characteristics with France (e.g. Belgium), up to a weight of 4 if they share none (e.g. the U.S.). Along a similar way of thinking, I propose other two measures in which I weight countries by their ranking in the Ease of Doing Business of the Work Bank: the first using the overall ranking and second using the specific ranking related to ease of trading across borders. Finally, I weight countries with the measure of network complexity proposed by Hausmann and Hidalgo (2011). In this case not only does coordination complexity include the number of export destination countries of a firm, but it also takes into account the complexity of the destination country itself. According to Hausmann and Hidalgo's definition, a country is more complex if it is more differentiated in the product space and at the same time it produces products that few other countries can make. They propose a theory where in order to produce a product a country needs to have all the necessary capabilities, hence very few countries will make products that require a lot of capabilities.

Table 2.6 displays the results. The regressors are not in logs but they have been standardized in order to compare them more easily. Capital, skill and HQ intensities are also included in the regressions but not displayed. Column (1) essentially corresponds to column (6) in Table 2.4. The gravity related measures in columns (2) and (3) give very similar results. The measures obtained weighting countries by their ease of doing business are still very robust but have a somewhat smaller magnitude. The intuition might come from the fact that these measures give a lot of weight to small developing countries. Even though fixed costs might be higher along some dimensions, they are actually lower along others. For instance the goods exported to those countries might be of lower quality, or less differentiated, hence less advertising is needed in order to penetrate the market.[33] In this respect it is interesting to note that the magnitude increases when countries are weighted according to the measure of complexity proposed by Hausmann and Hidalgo (2011). Complex countries are generally more advanced (their measure is correlated with income per capita) and the evidence seems to suggest that exporting to those countries is more difficult: more service inputs are required and firms tend to outsource them.

2.5.2.2 Selection into Exporting and other Robustness Checks

The results in the previous sections are obtained with variables in logs, hence they only include exporters and firms that do outsource at least some of their service inputs. On the one hand, this makes firms more comparable because they are likely to be more similar across different dimensions. On the other hand, selection might be an issue. Unfortunately the sample I have is not well suited to analyze this issue, for two main reasons. First, including mainly large firms, the survey is highly skewed towards exporters; hence non-exporters might not be fully representative for the population of firms. Second, the EAE dataset has the serious drawback of not distinguishing between zeros and missing values. Even though I try to solve this issue by imputing missing values (see Appendix 2.B.3), I cannot be entirely sure that firms reporting zero outsourced services are in fact firms that simply did not fill that section of the survey. Measurement error is therefore likely to be present and might affect the results. Despite these warnings, I re-obtain all the results without taking logs, and the exact same picture emerges. For instance Table 2.7 replicates the results of Table 2.4 without taking logs of the regressors but standardizing the variables to compare them more easily. Coordination complexity measured as the number of destination countries is still positive and highly significant in all specifications. Moreover the overall trade extensive margin (being an exporter) is also positive and highly significant in all specifications, while the trade intensive marking remains insignificant. A very similar picture holds true when also firms that do not outsource services are included.

Finally the baseline result is very robust across many specifications and controls. For instance the same picture holds true when purchased services are weighted by total sales instead of total costs. Moreover the EAE survey also contains a measure of outsourcing of non-core activities. The model does not differentiate inputs; therefore there is no 'a priori' clear distinction between a service or a non-service input, apart from the intuitive assumption that for manufacturing firms the importance of adaptation will be higher for the primary good inputs (hence, by Proposition 1, they will be more likely to be produced in-house compared to services). There are although manufacturing firms whose activity has almost completely shifted towards services, which have essentially become their core competencies (Nike and P&G are two leading examples). In this respect a measure of outsourcing of non-core activities is possibly even more in line with the

[33] An interesting avenue for future research is to test whether the form of fixed costs proposed by Arkolakis (2010) holds in the data. His interpretation in terms of marketing costs could be fairly easily tested with direct measures of advertising costs.

Table 2.7: PBS Outsourcing and Coordination Complexity with Non-Exporters

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Exporter	0.042***	0.035***	0.035***	0.035***	0.034***	0.034***	0.034***	0.034***	0.040***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.010)
NC		0.141***	0.141***	0.140***	0.135***	0.134***	0.132***	0.125***	0.125***
		(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.017)
Exports			0.003	0.003	-0.001	-0.001	0.002	-0.002	0.001
			(0.011)	(0.011)	(0.012)	(0.012)	(0.015)	(0.014)	(0.014)
Capital Intensity				-0.000	-0.002	-0.002	-0.002	-0.002	-0.001
				(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Skill Intensity					0.001	0.001	0.001	0.001	0.001
					(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
HQ Intensity						-0.005	-0.005	-0.005	-0.003
						(0.005)	(0.005)	(0.005)	(0.005)
Scale							0.024*	0.019	0.018
							(0.014)	(0.014)	(0.013)
Num imp. products								0.024**	0.029***
								(0.011)	(0.011)
Contract Intensity									0.003
									(0.008)
Observations	235,182	235,182	235,182	234,756	224,561	224,260	224,260	224,260	167,621
Number of firms	39,500	39,500	39,500	39,457	38,326	38,300	38,300	38,300	29,027
R-sq W	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of purchased services over total costs in logs. All other variables are standardized. See also notes in Table 2.4.

model. Managerial complexity, measured as the number of export destination countries, has again a positive and even stronger impact on outsourcing, when this alternative definition is considered.[34]

2.5.3 Outsourcing versus Internal Production

An alternative explanation for the rise in service outsourcing is that manufacturing firms are simply becoming more service oriented. An overall increase in the production of services might increase the need for service inputs, that in turn leads to more outsourcing. This could lead to a parallel increase of services produced internally and services purchased from the market. Controlling for internal production is therefore key in the analysis to rule out the possibility that service outsourcing is driven by manufacturing firms simply becoming "service firms". I propose six other measures of internal production and the base result is robust to all of them, as shown in Table 2.8. Column (1) displays the same regression of column (6) in Table 2.4, where I capture internal service production with the share of revenues generated by establishments classified in services (often the headquarters). In column (2), the definition of internal production of services is very similar but the headquarter share is computed in terms of employment and not revenues. The following two regressions employ similar definitions but instead of using the shares of all service establishments they only include the establishments classified in business service industries (for instance they exclude transportation, retail, wholesale etc...).[35] Columns

[34]Table 2.C.1 in the Appendix shows the results. An interpretation of this result is that most of these non-core activities are actually services given that only manufacturing firms are analyzed, the two measures of outsourcing could be therefore quite similar (see the definitions in Appendix 2.B.2).

[35]Whenever the shares in terms of employment are used (columns 2 and 4), year 2007 is dropped due to missing data.

(5) and (6) use respectively the share of total salaries and of total employment accounted in establishments classified as headquarters by the firm itself.

Table 2.8: Alternative Measures of Internal Production of Services

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
NC	0.086***	0.080***	0.086***	0.080***	0.086***	0.086***	0.086***
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
HQ Share (Rev)	-0.052						
	(0.037)						
HQ Share (Empl)		-0.005					
		(0.049)					
PBS Share (Rev)			-0.047				
			(0.369)				
PBS Share (Empl)				-0.210			
				(0.420)			
HQ Est. (Salaries)					0.053		
					(0.071)		
HQ Est. (Empl)						0.007	
						(0.075)	
Professionals Share (CS3)							0.043
							(0.084)
Observations	174,682	161,755	174,682	161,755	174,745	174,812	174,892
Number of firms	30,159	29,464	30,159	29,464	30,126	30,161	30,172
R-sq W	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of purchased services over total costs. All columns also include the following regressors: exports, capital and skill intensities. Year 2007 dropped in columns (2) and (4). See also notes in Table 2.4.

All these measures have the common problem that they do not account for services produced inside production establishments. This is a measurement issue is also present for the economy as a whole: when a manager sits in the back of a production site performing accounting, billing, marketing and other services, all this service activity goes completely undetected in industry data. This issue is probably not too worrisome in the present context since I mainly focus on exporters. In fact it is well known that exporters are larger, and large firms tend to have establishments dedicated to services, as reported by Young and Triplett (1996). In any case, I propose another measure of internal production that is not subject to this problem since it comes from occupation data. I use the share of workers classified as managers and professional occupations in total employment (column 7). The measure comes from the DADS dataset and it is a relatively good proxy for internal production of PBS services since these activities are mainly performed by professionals.[36]

Another possibility is that manufacturing firms are not becoming more service oriented but they "consume" relatively more services in order to export. This is of course at the heart of the mechanism under study. Firms need more services inputs to export to more destination countries and the overall increase in the need of services might exceed their expansion in terms of total revenues or total costs, hence the share of services would mechanically increase. Firms

[36] This is category 3 (CS3) in the DADS data: "cadres et professions intellectuelles supérieures". Ideally I would control for their share in the total employment bill but I do not have wages disaggregated at that level of detail.

Table 2.9: Total Service Production

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exporter	0.088***						
	(0.021)						
NC		0.105***	0.107***	0.108***	0.111***	0.118***	0.124***
		(0.012)	(0.013)	(0.013)	(0.013)	(0.014)	(0.015)
Exports			-0.001	-0.001	-0.000	-0.000	0.001
			(0.005)	(0.005)	(0.005)	(0.006)	(0.006)
Capital Intensity				0.032**	0.028**	0.014	0.010
				(0.014)	(0.014)	(0.015)	(0.015)
Scale					-0.039	-0.059**	-0.063**
					(0.025)	(0.026)	(0.027)
Num imp. products						0.023**	0.030***
						(0.009)	(0.010)
Contract Intensity							0.021***
							(0.007)
Observations	194,581	159,698	159,698	159,078	159,078	139,949	134,888
Number of firms	34,399	28,178	28,178	28,086	28,086	24,624	23,976
R-sq W	0.03	0.03	0.03	0.03	0.03	0.03	0.03
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of purchased services in total costs plus the share of professionals in total employment. All variables are in logs. Data are for period 1996-2007. Firm-clustered standard errors in parentheses; (*, **. ***) indicate 10, 5, and 1 percent significance levels.

Table 2.10: Outsourcing versus Internal Production

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exporter	0.088***						
	(0.021)						
NC		0.084***	0.079***	0.081***	0.064***	0.064***	0.067***
		(0.012)	(0.013)	(0.013)	(0.013)	(0.014)	(0.015)
Exports			0.004	0.003	-0.003	-0.003	-0.001
			(0.005)	(0.005)	(0.005)	(0.006)	(0.006)
Capital Intensity				0.016	0.036***	0.032**	0.030*
				(0.014)	(0.014)	(0.015)	(0.016)
Scale					0.213***	0.197***	0.194***
					(0.023)	(0.025)	(0.025)
Num imp. products						0.009	0.012
						(0.009)	(0.010)
Contract Intensity							0.012*
							(0.007)
Observations	194,581	159,698	159,698	159,078	159,078	139,949	134,888
Number of firms	34,399	28,178	28,178	28,086	28,086	24,624	23,976
R-sq W	0.01	0.02	0.02	0.02	0.02	0.02	0.02
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is ratio of the share of purchased services in total costs over the share of professionals in total employment. All variables are in logs. Firm-clustered standard errors in parentheses; (*, **. ***) indicate 10, 5, and 1 percent significance levels.

might decide to source them both from inside and outside the firm, and if the shift takes place at the exact the same pace, the boundary of the firm might not be an issue after all. At first sight, it looks like the overall need of services has in fact increased. Table 2.9 shows the sum of both purchased services and internal production of services, measured as the share of professionals in total employment. It is indeed strongly related to the number of destination countries. This fact implies that the impact of globalization on structural transformation might be even stronger than what the data shows. Many of these professionals are not employed in service establishments, hence they will be accounted within manufacturing when in fact they are producing services. The larger magnitude of the coefficient of interest with respect to the case of outsourcing only (Table 2.4) seems to point in that direction.

However the boundaries of the firms do matter. First of all, the magnitude for the overall production of services is larger than in the case of outsourcing only, but marginally. Thus internal production contributes much less then outsourcing, and exporting to more countries increases the outsourced share of services dis-proportionally more than internal production. This fact can be shown in two ways. First, in Table 2.10 I run a set or regressions where the dependent variable is the ratio of outsourcing over internal production (hence the log difference between the outsourcing share and the professionals share). The coefficients are positive and strongly significant, so an increase in the number of destination countries leads firms to increase outsourcing over internal production.

Second, it is evident from Table 2.11 that the magnitude of the effect for internal production alone (share of professionals, column 2) is significantly smaller than the one for outsourcing. The table also shows how other categories of occupations are related to the main variable of interest, these categories correspond to the hierarchical layers described by Caliendo et al. (2012). As expected the share of professionals (CS3, or Layer 2 in Caliendo et al. 2012) expands the most compared to all other internal layers, confirming that this is the occupational category most likely producing the specific services associated with exporting. Since the share of professionals is the one increasing the most with the number of destination countries, the exercise in Table 2.10 was the most demanding setting in which to test the ratio of the two. Table 2.11 also shows the ratio of outsourcing over internal production for all layers (columns 6-10, column 7 coincides with column 5 of Table 2.10). Outsourcing increases with the number of export destination countries disproportionally more than any measure of internal production. All in all, the results appear to be very robust to internal production of services. This evidence shows not only that business services are a fixed export cost component but also that firms tend to acquire these key inputs by outsourcing them to external providers, rather then producing them in-house.

An interesting point to note is that one of the empirical results that do not match with the theory in Caliendo et al. (2012) is related to the proportional expansion of higher layers with respect to lower layers. For firms that start exporting, the theory would predict that higher layers should expand more than lower layers, but the authors do not seem to find evidence for this theoretical prediction in the data. An explanation for this apparent puzzle is that the theory developed by Garicano (2000) and Caliendo and Rossi-Hansberg (2012) does not explicitly draw the boundary of the firm. So there is nothing that imposes that problem solvers, who have the knowledge to solve exceptional problems, should be employed directly by the firm.[37] My results strongly suggest that the expansion of higher layers come from across the boundary of the firm: firms outsource these high skill and infrequent services to external specialists. This strategy also allows firms to be more flexible, in fact they do not need to pay the fixed cost correspondent to the wage of the problem solver, but they can access his knowledge only when needed. In the Caliendo and Rossi-Hansberg's (2012) framework this could be seen a way to smooth the transition between different number of layers.

Table 2.11: Internal Hierarchies and Outsourcing

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Dependent Var.	CS2	CS3	CS4	CS5	CS6	OUT/CS2	OUT/CS3	OUT/CS4	OUT/CS5	OUT/CS6
NC	-0.015***	0.026***	0.023***	0.020***	-0.015***	0.078***	0.064***	0.062***	0.068***	0.097***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)	(0.015)	(0.013)	(0.012)	(0.012)	(0.012)
Export	-0.010***	0.001	0.000	-0.003*	0.002**	0.019***	-0.003	-0.002	0.003	-0.004
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
Capital Intensity	-0.022***	-0.003	-0.003	0.001	0.019***	0.061***	0.036***	0.037***	0.033**	0.013
	(0.005)	(0.004)	(0.004)	(0.004)	(0.003)	(0.016)	(0.014)	(0.012)	(0.014)	(0.014)
Scale	-0.695***	-0.134***	-0.089***	-0.098***	0.093***	0.756***	0.213***	0.153***	0.158***	-0.027
	(0.014)	(0.009)	(0.008)	(0.010)	(0.008)	(0.033)	(0.023)	(0.022)	(0.024)	(0.024)
Observations	104,420	171,963	186,270	178,793	188,624	96,898	159,078	171,851	165,286	173,584
Number of firms	22,640	29,141	30,899	30,256	31,318	21,722	28,086	29,719	29,126	30,098
R-sq W	0.18	0.13	0.04	0.09	0.05	0.04	0.02	0.01	0.02	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of different occupational categories in total employment (column 1-5) and the ratio of the share of purchased services in total costs over the share of the categories in total employment (column 6-10). All variables are in logs. See notes in Table 2.4.

2.5.4 Timing: Sunk versus Fixed Costs

An interesting question is whether service export costs are more of a fixed or sunk nature. In the former case the costs have to be paid every period, while in the latter they are paid once for all when a firm enters a new market. In their quantitative exercise Morales et al. (2011) find that fixed costs are somewhat larger. This contrasts with the results of Das et al. (2007), who find the opposite. Both papers obtain estimates of these costs from structural models, while here I can provide direct evidence on the nature of export costs (or at least part of them) from the data. To the best of my knowledge this is the first paper that addresses this issue directly, and I do so through a simple analysis on the timing of the incurrence of these costs. For each firm I define two dummy variables that identify: a) the year before entering a new market; and b) the year of entrance. The first variable is simply the lead of the second. If the fixed costs are mostly sunk in nature we would expect the costs to be paid before entering and being related to the number of countries that will be entered, regardless the number of countries that the firms is exiting.

This is precisely what happens, as Table 2.12 shows. Although not displayed, the regressions also includes exports, capital, skill and HQ intensities as regressors. Firms appear to increase the purchases of services mostly in the year before entrance, and not when they enter. Moreover I build two sets of variables depending whether a firm is also exiting countries or not: in the

[37] In fact Garicano and Rossi-Hansberg (2012b), using a very similar setting, talk more generally about "referral markets".

'gross' case I simply set the variable to one if a firm enters a country regardless of exit, in the 'net' case, on the other hand, I define entry only if the number of entered countries is higher the the number of exited countries (hence the variable is zero if a firm enters a country and exits another country at the same time). It turns out that the magnitude of the coefficient is considerably higher in the 'gross' case, hence the evidence points to costs that are sunk in nature since exit does not seem to matter much. This statement is confirmed when I analyze re-entry. I define two dummy variables as before but for cases in which firms re-entry countries in which they had already exported in the past. When they re-entry these countries without simultaneously entering other countries, the costs does not seem to move at all, if anything the share of outsourcing is decreasing.

Table 2.12: The Timing of Service Outsourcing

	(1)	(2)	(3)	(4)	(5)
NC	0.077***	0.086***	0.072***	0.074***	0.075***
	(0.013)	(0.014)	(0.012)	(0.013)	(0.013)
Export	-0.002	-0.003	-0.000	-0.001	-0.001
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Country Entry (t+1, gross)	0.029***				
	(0.009)				
Country Entry (t, gross)	-0.007				
	(0.010)				
Country Entry (t+1, net)		0.018***			
		(0.007)			
Country Entry (t, net)		-0.011			
		(0.007)			
Country Re-entry (t+1, only)			-0.016*		
			(0.009)		
Country Re-entry (t, only)			-0.011		
			(0.010)		
Continent Entry (t+1, gross)				0.006	
				(0.008)	
Continent Entry (t, gross)				-0.007	
				(0.008)	
Legal Sys. Entry (t+1, gross)					0.011
					(0.008)
Legal Sys. Entry (t, gross)					-0.004
					(0.008)
Observations	148,243	148,243	153,492	148,243	148,243
Number of firms	26,372	26,372	27,554	26,372	26,372
R-sq W	0.01	0.01	0.01	0.01	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of purchased services over total costs. All columns also include the following regressors: K/L, S/L, HQ Intensity, and scale. See also notes in Table 2.4.

It is informative to analyze cases in which costs are more likely to be variable in nature rather than fixed or sunk. For instance I look at employment outsourcing, that is, the use of temporary work from employment agencies. It is likely that firms use these services when they have some capacity constraints and decide to expand in a more flexible way, for instance to meet peaks of demand. Table 2.13 provide strong evidence in this direction. In the baseline regression (column 1) it is now the intensive margin of trade to be positive and significant and not the number of destination countries. This result supports the idea that temporary employment is not used to produce country-specific inputs. Looking at the timing, it is now the net entry that matters and the costs are mostly incurred in the year of entrance and not before, in full contrast with the previous results. Hence this type of labor input is likely to be related to variable costs, rather than country specific entry costs. Similar evidence, albeit somewhat weaker, applies for industrial outsourcing, hence outsourcing of goods and components rather than services. It is the trade intensive margin that is positive and strongly significant, even though also the number of countries is marginally significant as well.

Table 2.13: Employment and Industrial Outsourcing

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Var.	Empl	Empl	Empl	Empl	Ind	Ind	Ind	Ind
NC	-0.009	-0.011	-0.014	-0.010	0.016	0.021*	0.022*	0.019
	(0.010)	(0.011)	(0.013)	(0.011)	(0.010)	(0.012)	(0.013)	(0.012)
Export	0.010**	0.012**	0.012**	0.012**	0.028***	0.027***	0.027***	0.027***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)
Capital Intensity	0.056***	0.066***	0.066***	0.066***	0.014	0.022	0.022	0.022
	(0.011)	(0.012)	(0.012)	(0.012)	(0.012)	(0.015)	(0.015)	(0.015)
Skill Intensity	-0.051***	-0.042***	-0.042***	-0.042***	0.032***	0.025**	0.025**	0.025**
	(0.011)	(0.012)	(0.012)	(0.012)	(0.010)	(0.011)	(0.011)	(0.011)
HQ Intensity	-0.257***	-0.275***	-0.275***	-0.275***	-0.311***	-0.376***	-0.376***	-0.376***
	(0.038)	(0.044)	(0.044)	(0.044)	(0.046)	(0.056)	(0.056)	(0.057)
Scale	0.101***	0.108***	0.108***	0.108***	-0.133***	-0.152***	-0.153***	-0.152***
	(0.022)	(0.026)	(0.026)	(0.026)	(0.025)	(0.029)	(0.029)	(0.029)
Country Entry (t+1, gross)		-0.002				-0.008		
		(0.008)				(0.009)		
Country Entry (t, gross)		0.012				-0.013		
		(0.007)				(0.010)		
Country Entry (t+1, net)			0.008				0.007	
			(0.006)				(0.007)	
Country Entry (t, net)			0.017***				-0.005	
			(0.006)				(0.007)	
Continent Entry (t+1, net)				0.002				0.000
				(0.008)				(0.008)
Continent Entry (t, net)				0.011				-0.013*
				(0.008)				(0.008)
Observations	142,809	122,409	122,409	122,409	135,767	115,238	115,238	115,238
Number of firms	26,510	23,510	23,510	23,510	26,080	22,847	22,847	22,847
R-sq W	0.04	0.04	0.05	0.04	0.01	0.01	0.01	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: See notes in Table 2.6.

In order to open the black box of fixed/sunk export costs it is interesting to investigate which are the specific service inputs that a firm needs in order to export to new countries. I can answer this question drawing on the extra information contained in the ERSI survey. The ERSI survey provides more information about service outsourcing policies. In particular the survey asks firms whether they outsource any service among a list of 34 different types of quite detailed services. Most of them are classified in the Professional and Business Services industry, but there are also some transportation, financial and real estate services. Unfortunately the survey is available in 2005 only and for a smaller sample of firms. In fact it includes all firms with more than 250 employees but just a sample of smaller firms for a total of 4,745 manufacturing firms (after the data cleaning procedure described in Appendix 2.B.3), compared to 24,117 firms in the EAE in 2005.[38]

I run a set of separate Probit regressions for each service type to see which one is mostly related to the number of destination countries. The regression specification is as follows:

$$OUT^*_{is} = \beta_0 + \beta_1 NC_i + \boldsymbol{W'_{it}\beta_2} + \delta_j + \epsilon_i \tag{2.20}$$

where $\epsilon \sim N(0, \sigma^2)$, δ_j is an industry fixed effect, and OUT^*_{is} is a latent variable such that:

$$OUT_{is} = \begin{cases} 1 & \text{if } OUT^*_{is} > 0 \\ 0 & \text{otherwise} \end{cases} \tag{2.21}$$

Hence OUT_{is} is now a binary variable that takes value equal to one if the service input s is outsourced by firm i. The set of control variables $\boldsymbol{W_{it}}$ include total exports, capital, skill and HQ intensities, and scale. Table 2.14 reports the marginal effects at the mean for the number of destination countries and total exports. It turns out that advertising, R&D and IT consulting are the service inputs most highly related to an increase of export destination countries, while the intensive margin of trade again plays no role. These inputs were already contained in the outsourcing variable used in the previous section, so it is reassuring to see that the results are confirmed with this more detailed survey for a cross-section of firms. Other service inputs that are key in order to export to new countries are: insurance, legal services, translation and quite intuitively transportation and packaging. Results change very marginally if I use a logit or a linear probability model.

2.5.5 Endogeneity: Is Reverse Causality a Problem?

Despite controlling for unobserved time-invariant firm characteristics and industry specific time shocks, a full causal interpretation of the previous results might remain problematic. A potential concern is reverse causality: firms might decide to outsource for reasons unrelated with exporting, as a result of that they become more productive and this allows them to become exporters. Given the positive relationship between outsourcing and productivity often found in the literature, this possibility is certainly a concern. At the same time the previous evidence shows the systematic and very robust relationship between the number of destination countries and outsourcing. Hence the standard setting of Melitz (2003) in which there is one productivity threshold over which the firms start exporting would not be enough. I would need multiple thresholds and systematic association between outsourcing of any input (even unrelated to exporting), productivity growth and increase in the number of destination markets.

A first way of investigating this issue is looking at the channel directly. In particular I can test whether the same inputs outsourced to start exporting are also the ones that generate the highest productivity gains. The last two columns of Table 2.14 report the results, where I run

[38] The firms in the ERSI sample account for 50% of the total turnover of the French manufacturing sector in 2005 (aggregate data from Eurostat).

Table 2.14: The Anatomy of Service Inputs

Service Type		Out. 1			TFP	
Num.	Description	NC	Exports	Obs.	Out	Obs.
5.1	Advertising	0.089***	-0.014**	3,884	0.022	3,861
2.1	R&D	0.052***	0.006	3,890	0.029	3,867
1.1	IT Consulting	0.042***	0.002	3,897	0.030	3,873
4.9	Insurance	0.034***	-0.004	3,892	-0.009	3,869
4.4	Legal Services	0.029***	-0.001	3,893	-0.030	3,870
3.3	Packaging	0.027***	0.001	3,890	0.117**	3,867
3.1	Transportation	0.022***	-0.000	3,894	-0.131***	3,871
4.1	Translation	0.022**	0.019***	3,889	0.043**	3,866
1.2	IT Maintenance	0.021*	0.007	3,890	-0.029*	3,866
6.6	Refuse collection	0.017**	-0.002	3,892	-0.164***	3,869
4.2	Training	0.016*	0.001	3,894	-0.019	3,871
4.6	Recruitment	0.013	0.006	3,890	0.068***	3,867
4.3	Business Consulting	0.012	-0.006	3,889	0.001	3,866
4.10	Leasing	0.011	-0.009	3,890	-0.047	3,867
6.4	Security	0.011	0.002	3,893	-0.022	3,870
6.7	Real estate	0.010	-0.008	3,891	-0.015	3,868
3.5	Chartering	0.009	0.010*	3,892	0.029	3,869
4.11	Debt recovery	0.006	-0.002	3,888	-0.042**	3,865
5.2	After-sales Services	0.006	-0.003	3,880	0.002	3,857
4.5	Accounting	0.002	0.006	3,893	-0.030	3,870
4.8	Brokerage	0.002	0.002	3,869	0.041	3,864
6.3	Cleaning	0.002	0.002	3,893	0.006	3,870
7.2	Personal services	0.002	-0.001*	3,339	0.028	3,862
2.2	Technical studies	0.001	0.014**	3,873	0.086*	3,850
4.7	Temporary work	-0.000	0.005	3,893	0.109***	3,870
6.5	Sewage	0.000	0.001	3,890	-0.010	3,867
1.3	Data processing	-0.001	-0.002	3,890	0.020	3,866
6.1	Machinery Maint.	-0.002	0.001	3,893	-0.021	3,870
6.2	Buildings Maint.	-0.006	0.012**	3,892	0.030	3,869
7.1	Catering	-0.009	0.015**	3,884	0.011	3,862
6.8	Machinery Renting	-0.011	0.003	3,893	0.018	3,870
3.2	Warehousing	-0.018	0.038***	3,892	0.125***	3,869

Note: Data from the 2005 ERSI Survey. See Appendix 2.B for the precise definition of service types. The table reports the marginal effects at the mean for number of countries and total exports. The last two columns report the results of separate OLS regressions where the firm's TFP is regressed on the outsourcing binary variable for each service type. All regressions include K/L, S/L, HQ Intensity and scale as controls. All regressors are in logs, and only exporting firms are included. Standard errors are clustered at the NES36 (Nomenclature Economique de Synthèse - Niveau 2) industry level used to stratify the sample of firms; (*, **. ***) indicate 10, 5, and 1 percent significance levels. The results for TFP will be described in the next section.

separate OLS regressions for each service type to see whether outsourcing of that particular service has an impact on the firm's TFP. The table shows that there is no correspondence between the outsourcing of service inputs needed to export to new countries and an increase in the firm's TFP, the reverse causality channel does not seem to be there.

A second way is searching for some plausible instruments. In the ideal setting I would want an exogenous shock that makes exporting suddenly easier and more attractive; as a consequence firms start exporting (or will export to more destination countries), and I would like to observe whether they change their sourcing behavior at home, for services in particular. Not having such a shock for France over the past fifteen years, I proceed in two ways. First, I look at the export destination growth for the US. The growth in destination countries for US firms can be related to the export opportunities of French firms, but at the same time unrelated to their productivity gains due to outsourcing. Second, I look at shocks to demand that are plausibly exogenous to French firms. If the demand for French products increases globally, French firms will find exporting more attractive. The 'China shock' and the EU enlargement seem the obvious choice.

I instrument the increase in the number of destination countries with the plausibly exogenous increase in the number of country-product varieties exported by the US or imported by China and the new EU members. This approach is close in spirit to one of the instrumental variables proposed by Bloom et al. (2011), and resembles the "shift-share" IV strategy used in the labor literature (e.g. Card, 2001). Since I have firm and industry-year fixed effects in all regressions I need an instrument that varies at the firm level. To achieve that I exploit the information on the products exported by each firm. I define the number of destination markets of a firm in the initial year as the number of markets reached by the firm's most successful product. Then I keep the firm's product space fixed and calculate the increase in the number of destination markets for each product with the increase in the number of countries where the US exports that particular product, or the number of countries that supply that particular product to China or the new EU members. In all cases, in the country count I exclude France, the EU15 or the Eurozone countries (as of 2001). In any year, the number of destination countries is given by the number of destination markets of the most successful product in that year.

The instrument is constructed as follows:

$$IV\_NC_{i,t,x} = \max_{p \in P_{i,t_0}} \left\{ NC_{p,t_0} + \Delta NC^{x}_{p,t} \right\} \tag{2.22}$$

where p is a product exported by firm i, and P_{i,t_0} is the full set of products exported by that firm in 1996 or the first year in which it exports. NC_{p,t_0} is the number of countries where firm i exports product p in the initial year, and $\Delta NC^{X}_{p,t}$ is the increase between year t_0 and year t in the number of export destinations of the US or in import sources of China or of new EU members for product p (hence $x \in \{$*US-Exports, China-Imports, newEUmembers-Imports*$\}$). For each x, I construct three instruments depending on the countries that I exclude in computing ΔNC: France (exFRA), the Eurozone countries as of 2001 (exEZ12), and the EU15 countries (exEU15).[39] I impose $IV\_NC_{i,t,x} >= 0$, hence more precisely my instrument is defined as: $\hat{IV}\_NC_{i,t,x} = \max\{IV\_NC_{i,t,x}, 0\}$. Since I only include exporters in the regressions by taking logs of all variables, whenever the constructed instrument is zero in a given year, the firm will be dropped in that particular year (even if it is in fact an exporter). But results virtually do not change if I do not drop those firms by imposing $IV\_NC_{i,t} > 0$.

Table 2.15: Instrumental Variable Estimation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
IV	US-Exp exEU15	US-Exp exEZ12	US-Exp exFRA	China-Imp exEU15	China-Imp exEZ12	China-Imp exFRA	newEU-Imp exEU15	newEU-Imp exEZ12	newEU-Imp exFRA
NC	0.549***	0.557***	0.565***	0.475***	0.452***	0.505***	0.390***	0.376***	0.398***
	(0.162)	(0.158)	(0.160)	(0.153)	(0.152)	(0.144)	(0.131)	(0.131)	(0.129)
Export	-0.100***	-0.103***	-0.103***	-0.085**	-0.080**	-0.092***	-0.065**	-0.062**	-0.066**
	(0.035)	(0.034)	(0.034)	(0.033)	(0.033)	(0.031)	(0.028)	(0.028)	(0.028)
Capital Intensity	0.035***	0.034***	0.034***	0.031**	0.032**	0.032**	0.032***	0.033***	0.033***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.012)
Skill Intensity	0.039***	0.039***	0.037***	0.043***	0.043***	0.041***	0.043***	0.043***	0.043***
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
HQ Intensity	-0.074*	-0.076**	-0.077**	-0.076**	-0.070*	-0.081**	-0.074**	-0.073*	-0.075**
	(0.039)	(0.039)	(0.039)	(0.038)	(0.038)	(0.038)	(0.038)	(0.038)	(0.038)
Scale	-0.002	-0.002	-0.005	0.003	0.009	-0.001	0.023	0.026	0.021
	(0.038)	(0.038)	(0.038)	(0.037)	(0.037)	(0.035)	(0.034)	(0.034)	(0.034)
Observations	160,887	160,799	160,917	164,546	164,638	164,960	166,092	165,659	165,764
Number of firms	24,061	24,058	24,055	24,386	24,383	24,412	24,539	24,514	24,523
F-stat	320.8	336.2	330.7	348.7	357.1	402.3	488.9	483.2	502.5
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: See notes in Table 2.6. The F-stat is the Kleibergen and Paap (2006) Wald rk F statistic provided by the Baum et al.'s (2010) xtivreg2 Stata command.

Table 2.15 shows the results of the exercise. Coordination complexity measured as the number of destination countries is again positive and very significant across all specifications. It also reassuring to see that the magnitude of the effect does not change much depending on the instrument used, and all of them are very strong as the F statistics from the first stage show. In Table 2.16, I perform a more demanding exercise where I only consider the change (hence the growth since all variable are in logs) between three years before and three years after the EU enlargement that took place in 2004. The model is therefore estimated in first differences between 2001 and 2007. The effect is still present and the magnitude is again very similar and robust to all the controls used in the regression.

Interestingly the previous OLS regressions were strongly underestimating the effect of interest. There could be several explanations for this result. But one in particular comes straight from the model. Firms with managers with higher ability will tend to produce more in-house, everything else constant (Proposition 3). From the trade literature we know that exporters, and in particular firms that export to multiple destinations, are at the very top of the firm distribution, hence they are much more likely to employ better managers and to produce internally at a lower cost. For all these reasons, firms that export to multiple destinations are more likely to produce inputs in-house, everything else constant. If they had to export to many destinations, less productive firms would tend to outsource a much higher proportion of their inputs, because

[39] The Eurozone countries in 2001 were: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, and Spain. The EU15 group include the previous countries and Denmark, Sweden, and the United Kingdom.

Table 2.16: IV Estimation - New EU member countries: 2001-2007 change (growth)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ IV	newEU exEU15	newEU exEU15	newEU exEU15	newEU exEU15	newEU exEU15	newEU exEU15	newEU exEZ12	newEU exFRA
Δ NC	0.296**	0.413**	0.410**	0.489***	0.449**	0.432**	0.450**	0.475***
	(0.121)	(0.191)	(0.200)	(0.185)	(0.191)	(0.199)	(0.192)	(0.183)
Δ Export		-0.074	-0.076	-0.094*	-0.084	-0.084	-0.088*	-0.095*
		(0.052)	(0.054)	(0.051)	(0.052)	(0.052)	(0.051)	(0.049)
Δ Capital Intensity			0.056	0.073*	0.071*	0.076**	0.076**	0.077**
			(0.037)	(0.037)	(0.037)	(0.037)	(0.037)	(0.037)
Δ Skill Intensity				0.080**	0.075*	0.080**	0.080**	0.079*
				(0.039)	(0.039)	(0.040)	(0.041)	(0.041)
Δ HQ Intensity					0.018	0.015	0.013	0.010
					(0.063)	(0.062)	(0.060)	(0.060)
Δ Scale						0.060	0.057	0.053
						(0.057)	(0.057)	(0.058)
Observations	9,234	9,234	9,180	8,494	8,381	8,381	8,380	8,383
F-stat	165.3	164.8	172.3	126.7	133.2	108.8	110.1	107.4

Note: The dependent variable is the growth of the share of outsourcing in total costs (log-change) between 2001 and 2007. All controls are also in log-changes. Standard errors are clustered at the 2 digit industry level. The F-stat is the Kleibergen and Paap (2006) Wald rk F statistic.

they would not be able to compete with external specialized suppliers. In reality these firms export to much fewer markets, if any at all. Hence if the number of destination countries could be randomly assigned to firms, the effect on outsourcing would be much stronger.

Considering the coefficients of the previous two tables, the average increase in the number of destination countries explains between 66% and 100% of the increase in the share of outsourced services over the period. The new channel between globalization and structural transformation that I propose is not only present but it is also quantitatively very significant. The causal effect of globalization essentially explains almost all of the increase in domestic service outsourcing observed in the sample.

One reason why the overall effect might be partially overstated is that the average effect might not be representative for the entire distribution of firms, due to the presence of non-linearities, as the next section will show.

2.6 Non-linear Effects: Evidence on Proposition 4

A further interesting question to answer is whether the effect of coordination complexity on the share of outsourcing exhibits a non-linear behavior. Proposition 4 predicts that the relationship should be concave, that is, the increase in outsourcing should flatten when the number of destination countries is large. This is precisely what happens, as Figure 2.7 shows.[40]

Multiple interpretations can be put forward to explain this fact. In the model, when the number of overall inputs goes up, the probability that the manager needs to understand the

[40] This effect is also confirmed by the negative sign on the square of the number of countries. The results of the regression are not reported but available on request.

input condition for each one of them becomes smaller and smaller. Setting up a communication code for such a high number of very rare events (all equally likely) is going to be very costly. So the manager increases the number of inputs internally produced at a much lower speed compared to the increase in the number of overall inputs needed. At some point the number of in-house produced inputs hardly increases and if the denominator keeps rising the overall share will become smaller and smaller till converging to zero. In the model the slope essentially goes to zero only when also the share of internally produced inputs goes to zero (even though with a lower order). In the data the slope is zero (even though not significant) when the share of internal production is still positive. The reason is of course due to the fact that inputs are homogeneous in the model while they are not in the data. It might be that the firm is outsourcing all of the service inputs needed to export but it is still of course producing in-house all the other core activities. One way to capture this effect in the model would be to have some heterogeneity in the importance of adaptation (e.g.: different probabilities of adaptation). In such case, the core activities would intuitively have a higher probability of adaptation and hence would be more likely to be produced in-house.

A different explanation is that the distribution of events might change. If the service inputs are not truly country-specific, but some of them are the same across countries in a differential way (e.g.: translation is shared across the countries with the same language, while transportation services are shared across all countries), then it is possible that some inputs are recurring with a higher and higher probability. In this situation the firm needs the input in different proportions and the uniform distribution of events would not apply anymore. Since some inputs become more and more frequent, the firm will find it optimal to produce them in-house because a code designed to communicate those events is quite cheap. The firm essentially specializes in the production of export-related services as it exports to more and more countries, and the share of service outsourcing might even decrease.

As Figures 2.7 shows, the data are very noisy for firms that export to a very high number of countries. In fact it is not possible to sign the slope, because the confidence interval allows for both positive and negative slopes. So it is hard to disentangle the two stories in the data.

2.7 Conclusions

By advancing the complexity of coordination, intrinsic to the managerial activity, as one of the main determinants of integration costs, this paper offers a better understanding of how the boundary of the firm are determined in presence of multiple inputs for which asset ownership is not very important, like many services. I have looked at one possible driver of coordination complexity: the internationalization decision of the firm. Exporting to more destination countries implies that more inputs are needed, and the higher number of inputs increases coordination costs making market transactions more appealing. I find new systematic evidence about domestic service outsourcing: an increase in the number of export destination countries has a strong positive effect on the share of purchased business services in total costs. This result establishes a new causal effect of globalization on structural transformation, which is quantitatively very significant. In fact the IV estimates, based on plausibly exogenous demand shifts, show that the average increase in the number of destination countries explains almost all of the increase in domestic service outsourcing observed in the sample.

Finally the paper makes a significant step forward in understanding the nature of export costs. Firms need to access a variety of specialized services to be able to export; often they do not have the capabilities to produce these inputs in-house so they have to rely on external suppliers. And the effect is stronger, the higher the number of markets that need to be reached. A flourishing and productive business services industry is therefore a key ingredient for a country export success, and its competitiveness on the world markets.

Figures

Figure 2.3: Service Sector Growth in France

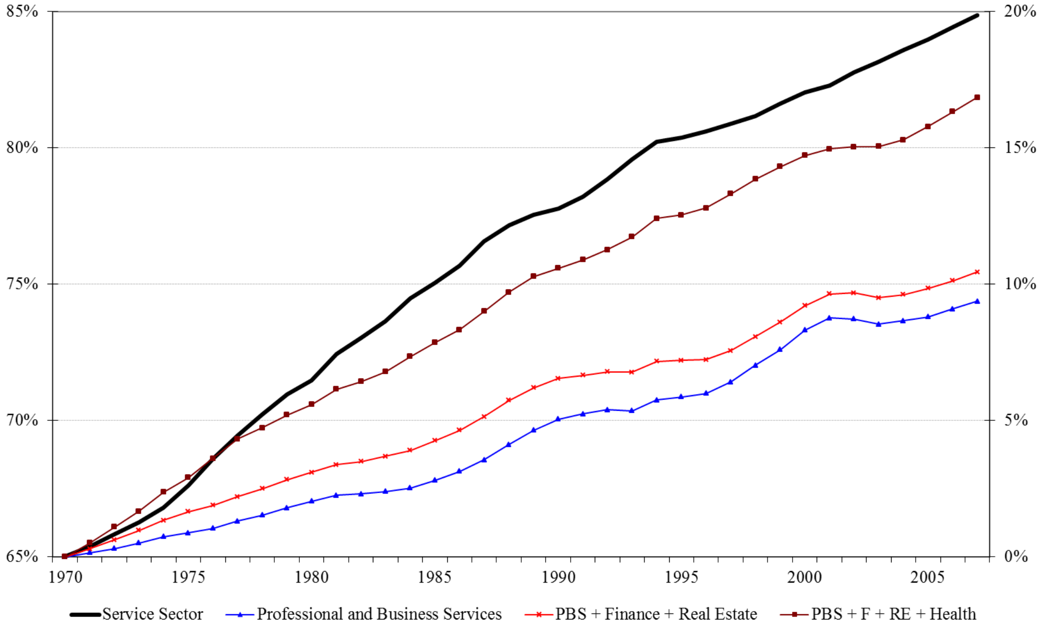

Source: KLEMS Dataset.
Note: The left-hand side axis displays the absolute share of the entire service sector (thick black line) in terms of total employment. The right-hand side axis applies to all series and displays the change in percentage points of total employment. The triangle marked line represents the percentage point change of Professional and Business Services (PBS); the cross marked line for the combined sector PBS, Finance and Real Estate; analogously the square marked line for the combined sector PBS, Finance, Real Estate and Health Care.

Figure 2.4: The Influence of PBS on the French Economy

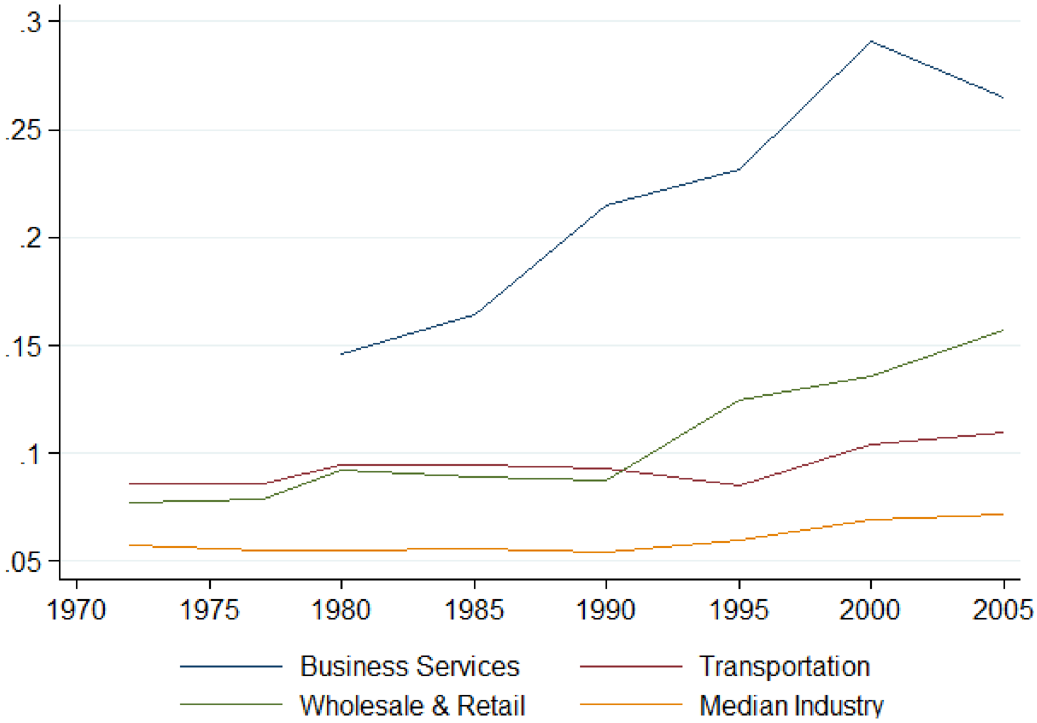

Source: KLEMS Dataset.

Figure 2.5: Share of Service Outsourcing and Relative Factor Intensities

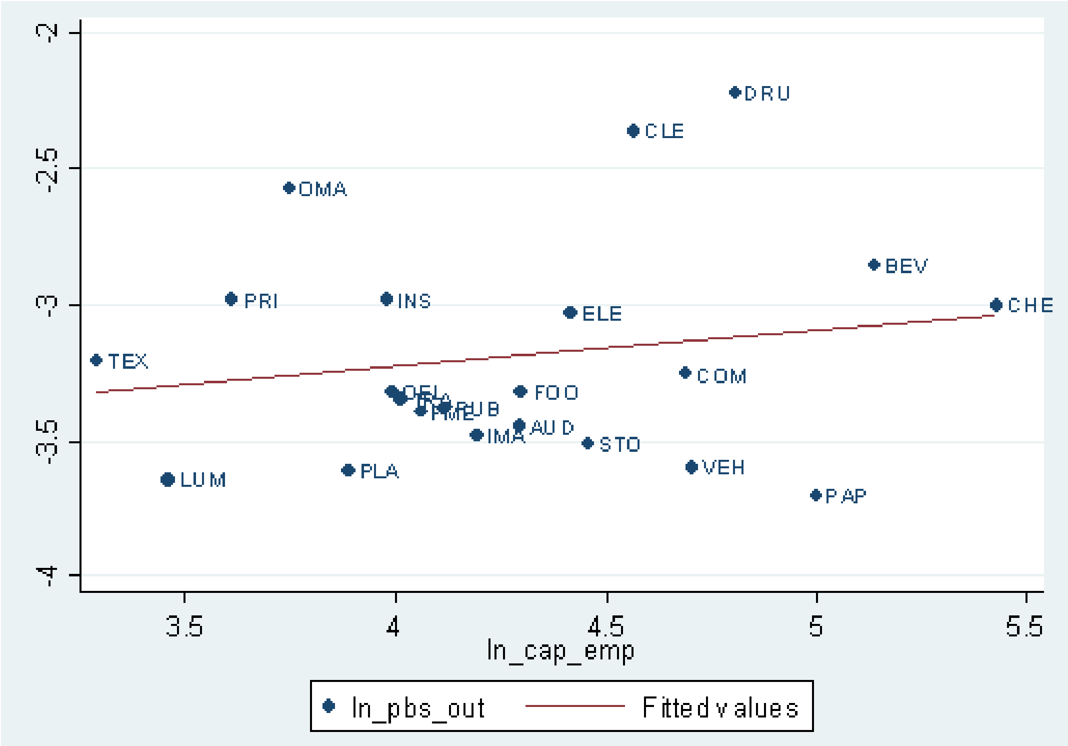

Source: BEA Bechmark Industry Accounts, NBER-CES Manufacturing Industry Database.
Note: All data are for 1992. The set of industries is defined as in Antràs (2003) apart from "other chemical products" that is combined with "Industrial chemicals and synthetics"; concordance tables are available on request.

Figure 2.6: Share of Service Outsourcing over Number of Destination Countries

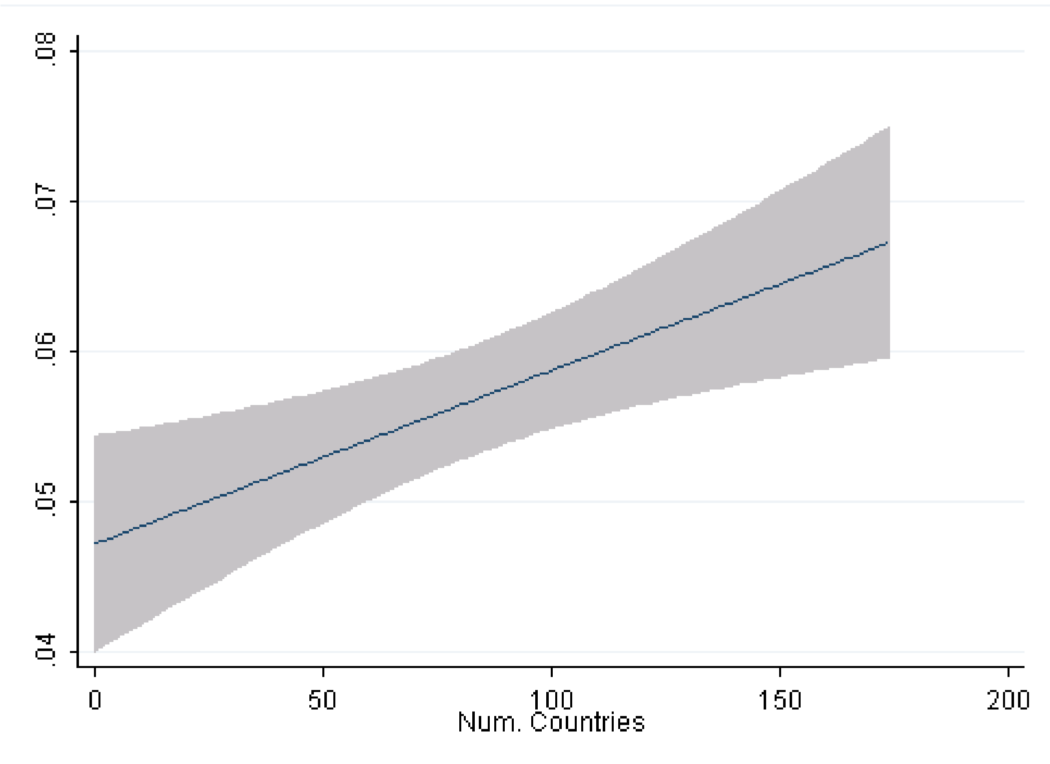

Source: The Enquête annuelle d'Entreprise (EAE) and French Customs data, 1997-2007.

Figure 2.7: Share of Service Outsourcing over Number of Destination Countries

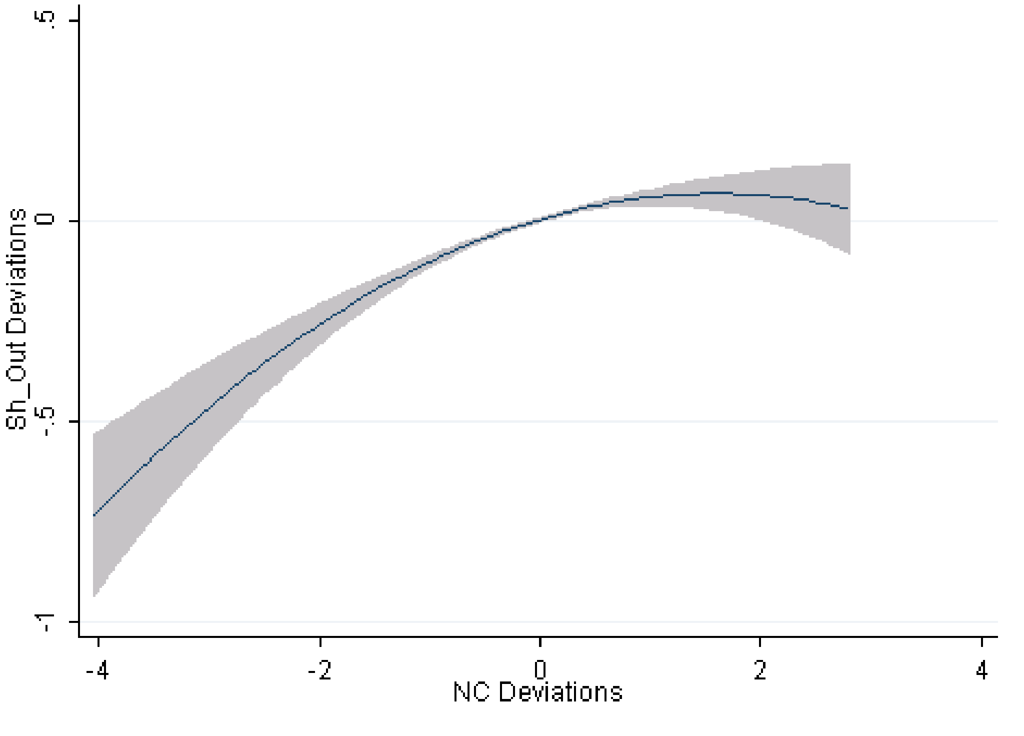

Source: The Enquête annuelle d'Entreprise (EAE) and French Customs data, 1997-2007.
Note: The fitted relationship corresponds to the best fitting quadratic functional form. The shaded area indicates 95% confidence intervals. The y-axis and the x-axis depict the residuals of two regressions of the log share of outsourcing or the log of the number of countries on firm and year-by-industry fixed effects. A very similar picture is obtained when also all other controls are added to the regressions.

Appendix

2.A Extensions to the Model

2.A.1 Enforceable Contracts

This appendix investigates what happens when contracts are enforceable by an external court. The reason why the firm decides to outsource is precisely not to pay the monitoring costs. Assuming that there exists an external court that can do that for free implies assuming the problem the away, and hence this setting is not fully in-line with the rest of the paper. Still, it is reassuring that all the main effects of interest are also present in this setting, and this extension offers interesting and intuitive results on the effect of institutional quality.

If an external court can fully enforce the contract (at no cost), the firm can specify a full contract with price and action even in the case of outsourcing. The problem is that the firm is not investing in the technology to communicate with the external supplier, so the manager will only know the expected value and not the actual realization for the input conditions that are outsourced. The manager will then solve the following problem:

$$\min_{\{a_i^v\},\{a_j^o\}} Nw_0 + \int_0^t (a_i^v - \theta_i)^2 di + \delta \int_0^t (a_i^v - \hat{\theta}^m)^2 di + E\left[\int_t^N (a_j^o - \theta_j)^2 dj + \delta \int_t^N (a_j^o - \hat{\theta}^m)^2 dj\right] \tag{2.A.1}$$

The optimal internal action does not change and it is easy to show that the optimal action for the generic outsourced input j is:

$$E[a_j^{o*}(\theta_j, \hat{\theta}^m)] = \hat{a}_j^{o*}(\hat{\theta}_j, \hat{\theta}^m) = \frac{1}{1+\delta}\hat{\theta}_j + \frac{\delta}{1+\delta}\hat{\theta}^m \tag{2.A.2}$$

At time zero, the manager will then sign a contract with the external supplier that specifies the tuple $\{P_j, \hat{a}_j^{o*}\}$, where $P_j = w_0 + (\hat{a}_i^{o*} - \theta_j)^2$. The payment is again a fixed price but the court will check that the external supplier will not deviate ex-post and will enforce the action $\hat{a}_j^{o*}$. What the firm is achieving is what Dessein and Santos (2006) define ex-ante and rigid coordination. In fact the manager does not give any flexibility to the external supplier but asks him to implement a specific action that is at least good on average, and hence can save part of the coordination costs ex-post.

It is easy to show that the total expected costs will be lower and equal to:

$$E[C]^{enf} = Nw_0 + \left[\frac{\delta}{1+\delta}t + \delta(N-t)\right]\sigma^2 + N\frac{\delta}{1+\delta}r^2 + M(t,N,K) \tag{2.A.3}$$

This is intuitive since the firm can achieve some degree of coordination despite avoiding the monitoring costs thanks to contract enforceability by court. Moreover the optimal share of inputs internally produced becomes:

$$\frac{t^{*enf}}{N} = \sqrt{\frac{K\sigma^2}{(1+\delta)N}} \tag{2.A.4}$$

This share is lower than the optimal share of the baseline case as long as $\delta > \delta^* = \sigma/\sqrt{\sigma + r}$. Hence, if adaptation is important enough, the share of outsourced inputs will be higher. This implies that firms in countries with better contracting institutions will be in general better off and will outsource a higher share of their inputs.

2.A.2 Adapting to the Average Action

This appendix solves a more general case, where the firm coordination costs are computed with respect to the average action taken by the firm for the inputs internally produced. To show the full set of implications, I start from the problem in a discrete setting and then see what happens when I move to the continuous case. The problem of the manager, once the input conditions have been realized, becomes:

$$\min_{\{a_i^v\}} Nw_0 + \sum_{i\in T}(a_i^v - \theta_i)^2 + \delta E\left[\sum_{i\in T}(a_i^v - \bar{a})^2 + \sum_{j\notin T}(a_j^o - \bar{a})^2\right] \tag{2.A.5}$$

where $\bar{a} = \frac{1}{N}\sum_0^N a_i$, and T is the set of inputs produced in-house. Solving for the optimal internal action is much more tedious and requires inverting a t-by-t matrix, but it can be shown that:

$$a_i^{v*}(\{\theta_i\},\{\hat{\theta}_i\},\{\hat{\theta}_j\}) = \hat{a}^v + \frac{1}{1+\delta}(\theta_i - \hat{\theta}_i) + \frac{\delta}{1+\delta}\frac{1}{N+\delta(N-t)}\sum_{i\in T}(\theta_i - \hat{\theta}_i) \tag{2.A.6}$$

where $\hat{a}^v = \frac{1}{1+\delta}\hat{\theta}_i + \frac{\delta}{1+\delta}\frac{1}{N+\delta(N-t)}\sum_{i\in T}\hat{\theta}_i + \delta\frac{1}{N+\delta(N-t)}\sum_{j\notin T}\hat{\theta}_j$. Hence the actions are fully interdependent: the optimal action for input i depends on all the average actions, plus the realizations of all the internal input conditions. The realization of the local input condition i still gets a higher weight compared to all others input conditions, but in order to internalize all externalities the manager moves away from that particular input condition to get closer to all other inputs. The optimal action is therefore a weighted average of all input conditions.

It is then possible to show that the expected costs become:

$$E[C] = Nw_0 + \left[\frac{\delta}{1+\delta}\frac{N+\delta(N-t)-1}{N+\delta(N-t)}t + \delta\frac{N-1}{N}(N-t)\right]\sigma^2 + \\ + \frac{\delta}{1+\delta}\sum_{i\in T}(\hat{\theta}_i - \tilde{\hat{\theta}})^2 + \delta\sum_{j\notin T}(\hat{\theta}_j - \tilde{\hat{\theta}})^2 + M(t,N,K) \tag{2.A.7}$$

where $\tilde{\hat{\theta}}$ is a weighted average of the means of the input conditions and is defined as follows: $\tilde{\hat{\theta}} = \frac{1}{N+\delta(N-t)}\sum_{i\in T}\hat{\theta}_i + \frac{1+\delta}{N+\delta(N-t)}\sum_{j\notin T}\hat{\theta}_j$. Further assuming that all input conditions have the same mean, the expression becomes:

$$E[C] = Nw_0 + \left[\frac{\delta}{1+\delta}\frac{N+\delta(N-t)-1}{N+\delta(N-t)}t + \delta\frac{N-1}{N}(N-t)\right]\sigma^2 + M(t,N,K) \tag{2.A.8}$$

This expression generalizes the expected costs of the baseline model. It is clear that the returns of integration are not constant anymore but depend on both t and N. It is possible to show that t and N are complementary but become substitutes if t is large. This implies that the advantage of in-sourcing diminishes when N grows large. The reason is that when the number of inputs is very large the manager cannot really achieve much by coordinating all the inputs in-house because the dispersion is too high.

Interestingly, when the model is written in the continuous case, the expected costs take the same form as in the baseline model:

$$E[C] = Nw_0 + \left[\frac{\delta}{1+\delta}t + \delta(N-t)\right](\sigma^2 + r^2) + M(t,N,K) \tag{2.A.9}$$

It is possible to formally show this result by re-solving the entire problem in continuum, or, more simply, by extending equation (2.A.8) in the limit. It is clear that the terms containing discrete elements tend to 1 (e.g. $(N-1)/N$). The intuition is that when the number of inputs grows very large, it is essentially not possible to internalize all the small externalities on the mean action, or more precisely, they do not matter in expectation. Mathematically all the interactions become quantities of a lower order when N grows large. Since the expected costs are the same, all the results in the main body of the paper apply.

2.B Data

2.B.1 Data Description

The industry level data come from the KLEMS database, while input-output data come from the OECD STAN database. Professional and Business Services include (Nace Rev 1 industry in parenthesis):

- Renting of machinery and equipment (71);
- Computer and related activities (72);
- Research and development (73);
- Other business activities (74), which include: legal, accounting, book-keeping and auditing activities; tax consultancy; market research and public opinion polling; business and management consultancy; architectural, engineering and other technical activities; advertising; labor recruitment and provision of personnel; investigation and security activities; industrial cleaning; miscellaneous business activities n.e.c.

The French micro-data come from the following four main data sources:

1. The Enquête annuelle d'Entreprise (EAE) that collects balance sheet data on all French firms with more than 20 employees and a sample of smaller firms;

2. The Déclaration annuelle de données sociales (DADS) that collects employment data on all firms with paid employees; the data used are aggregated at the establishment level;
3. Transaction level import-export data come from the French Customs. These data have been used among others by Eaton et al. (2004);
4. Finally service outsourcing data contained in the EAE are integrated with the Enquête Recours aux Services par l'Industrie (ERSI), a survey of firms with more than 20 employees and the census of firms with more than 250 employees that collects detailed information about service outsourcing policies for the year 2005. The total response rate of the survey is 85% and is well-balanced across industries and firm sizes.

Data for the gravity variables are provided by Mayer and Zignago (2011).

2.B.2 Variable Definitions

2.B.2.1 The Enquête annuelle d'Entreprise and DADS

Exporter: the variable takes the value of 1 if the firm reports positive exports in the Custom data, 0 otherwise (including the firms that do not appear in the Custom data).

Capital Intensity (K/L): ratio of the total capital stock to total employment, where the capital stock is measured as the total of tangible capital assets at end of year (I150) and total employment is the total number of full time equivalent employees (E101).

Headquarters intensity: ratio of workers employed in branches that produce services (Nace codes from 50 to 93) to total employment. It is a measure of internal production of services.

Contract intensity: the variable is constructed using the information about firms' imports. The firm-level contract intensity is therefore a weighted average of the contract intensity of all firm imports, where the measure of contract intensity is taken from Rauch (1999), and the weights are the shares of each product in the total firm imports. An imported good is considered as contract intense if it is neither sold on an organized exchange nor reference priced, I use the Rauch's (1999) 'Liberal' classification as in Nunn (2007) and Corcos et al. (2013).

PBS Outsourcing Share: in the baseline case it is defined as the sum of purchases of studies (D321), purchases of IT services (D329), and advertising (D360) over either Total Costs or Turnover (R310).

Professionals Share: it is the share of workers classified as managers and professional occupations (cadres et professions intellectuelles supérieures) in total employment. It comes from DADS and it is another proxy for internal production of PBS services, given that they are mainly produced by professionals.

Scale: total number of full time equivalent employees (yearly average, E101).

Skill Intensity (S/L): ratio of skilled workers to unskilled workers (from DADS). The number of skilled workers is the sum over all establishments of non-secondary jobs at the end of the year for the following categories: chief executives (chefs d'entreprises salariés), managers and professional occupations (cadres et professions intellectuelles supérieures), intermediate professions and technicians (professions intermédiaires). Unskilled workers include the following categories:

sales and administrative occupations (employés), qualified and unqualified operators and laborers (ouvriers). All of the previous categories include ordinary employment only and exclude for instance interns and apprentices.

Value Added over Sales (VA/Sales): ratio of value added to turnover (R310). Value added is defined as turnover minus purchases of goods (R210) and purchases of raw materials (R212). In the baseline definition I do not use other purchases and charges (R214), and other charges (R222) because they also include some labor costs. The former contains charges for external personnel (D350), while the latter board of directors' fees.

TFP: It is computed using the Levinsohn and Petrin's (2003) methodology. The coefficient of a Cobb-Douglas value-added production function are estimated at the 3 digit NACE industry level using raw materials (R212) as the proxy for the productivity shock. TFP at the firm level is then calculated as a residual between the actual and predicted Value Added using the estimated coefficient.

Total costs: it is used to calculate the outsourcing shares. It is the sum of purchases of goods (R210), purchases of raw materials (R212), other purchases and charges (R214), total labor costs (R216), social contributions (R217), and other charges (R222).

2.B.2.2 ERSI

The ERSI survey contains information about 34 types of services; in particular for each service type it provides a binary variable equal to one if the service is outsourced by the firm. Hence the OUT_{is} variable corresponds to the B* variables contained in the survey. I use the revised version of the variables, adjusted to remove internal inconsistencies. The service types are:

1. ICT Services
   - 1.1: IT consulting
   - 1.2: Software and IT third party maintenance
   - 1.3: Data processing and IT management
   - 1.4: Telecommunications
2. R&D and Professional Services
   - 2.1: Research and development
   - 2.2: Architecture, engineering and technical studies
3. Transportation Services and Logistics
   - 3.1: Railways, air, water and land transport
   - 3.2: Handling and warehousing
   - 3.3: Packaging
   - 3.4: Courier and post
   - 3.5: Chartering and international transport
4. Administrative Services
   - 4.1: Secretariat, translation and interpreting
   - 4.2: Vocational training
   - 4.3: Business and management consultancy
   - 4.4: Legal services

4.5: Accounting, book-keeping and auditing
4.6: Labour recruitment and provision of permanent personnel
4.7: Temporary work
4.8: Securities broking and fund management
4.9: Insurance and other financial services
4.10: Leasing
4.11: Invoicing/billing and debt recovery

5. Commercial Services
   5.1: Advertising, marketing and communication
   5.2: After-sales services

6. Maintenance and General Services
   6.1: Car, equipment and machinery maintenance
   6.2: Buildings maintenance
   6.3: Cleaning
   6.4: Investigation and security activities
   6.5: Sewage and sanitation
   6.6: Refuse collection, treatment and recycling
   6.7: Real estate
   6.8: Renting of machinery, car and transport equipment

7. Personnel Services
   7.1: Restaurants, canteens and catering
   7.2: Day care, nurseries and personal services

2.B.2.3 Robustness Checks

Scale_2: sum of non-secondary jobs at the end of the year over all establishments from DADS (EFF_3112_ET).

K/L_2: capital intensity where capital is measured as before but total employment is taken from DADS (EFF_3112_ET).

Out_2: compared to the baseline case it adds non-capital expenditures on software purchases (D511). It is again computed as a share of either Total Costs or Turnover (R310).

Out_3: compared to Out_2 it adds capital expenditures on software purchases (I461) and investment in R&D (I122). It is again computed as a share of either Total Costs or Turnover (R310). The investment in R&D corresponds to expenses of the firm due to the acquisition, the creation, the provision, or the transfer of R&D in the current year. It is not possible to rule out the possibility that some of these expenses are actually incurred within the firm, hence this is probably the least reliable measure of outsourcing.

Out_1b, Out_2b, Out_3b: compared to their respective cases they include outsourcing of non-core activities (D323) instead of purchases of studies (D321). Outsourcing of non-core activities corresponds to item 611 of the French national accounting code (Plan Comptable Général - PCG), which is defined as the outsourcing of tasks not specifically related to the core business of the firm and not already counted in item 604 of the PCG (D321).

S/L_2: skill intensity calculated as the ratio of skilled workers to total workers.

VA_2: value added defined as the sum of turnover (R310) and other goods sold (R315) minus purchases of goods (R210), purchases of raw materials (R212), and other purchases and charges (R214). Note that other purchases and charges contain the cost of external personnel (payments to employment agencies).

2.B.3 Data Cleaning

All variables from EAE before 2001 and salary from DADS before 2000 are transformed into euro. Unfortunately there are no missing values in the database and all variables are zeros even when they are clearly missing. So I set the relevant variables to missing in the following cases:

- If all balance sheet variables are zeros (E* R* D* I* S*);
- If all income statement variables are zeros (R1* R2* R3* R40* D* S001);
- If all cost variables are zeros (R2*);
- If all employment variables are zeros (E* S003 D350 D351 D352 - after having performed the adjustments described below);
- If employment is zero (E101) but total labor costs are positive (R216);
- If all intangible investment variables are zeros;
- If capital stock is zero (I150);
- If purchases of studies (D321) and purchases of materials (D322) are zeros but the variable containing their sum (D314) is positive;
- If all outsourcing and external charges are zeros (D3* D5*);

The following adjustments are also performed: *Capital* (I150): whenever possible, I obtain the end of the year capital stock from the stock at the beginning of the year by adding acquisition and revaluations and subtracting decumulation and disposals.

Exporter and other trade variables: I set them to missing (from zero) if a firm is reporting positive exports in EAE but exports are missing in the Custom data. This is mainly true for small exporters within the EU, who are not required to fill Customs data if the total value of annual exports is below 100k€[41]. At the same time there are also some cases of large exporters, this could be due for instance to confidential trade.

Other purchases and charges (R214): whenever it is zero or too small I take the sum of its components, which, according to the French accounting rules (Plan Comptable Général), are: outsourcing of non-core activities (D323), payments for leasing (D330), salaries to external employees (D350), advertising (D360).

Employment (E101): employment is measured as the total number of full time equivalent employees (annual average). Whenever possible, I replace the zeros with the sum of the annual average employment over all branches (S003), or with employment at the end of the year (E200), or with with the sum of the annual average employment over all establishments (V001), or, finally, with the sum of non-secondary jobs at the end of the year over all establishments from DADS (EFF_3112_ET). I use employment at the end of the year from DADS instead of the

[41]The reporting threshold is 100k€in the 2001-2005 period, 150k€in 2006 and 2007, and 38k€(250k Francs) before 2001

annual average employment (EFF_MOY_ET) because the latter is not available before 2002; non-secondary jobs (postes non annexes) exclude secondary jobs that last or are paid too little (see INSEE, 2013). When I use EFF_3112_ET as a robustness for capital intensity, I replace the missing and zeros with E101 to keep the same sample size.

Headquarters intensity: when calculated in terms of labor shares, it is set to missing in 2007 because very few firms report employment by branch in that year.

Outsourcing shares: firms are dropped whenever any of the outsourcing shares (in terms of turnover or total costs) exceed one.

Purchases of goods (R210), *Purchases of raw materials* (R212): they are set to missing if negative (only few cases in 1996).

Purchases of Studies (D321): the sum of purchases of studies and purchases of materials (D322) is contained in the outsourcing of activities related to the core business (D314). I calculate the average share of purchases of studies in the total at the two digit Nace industry level. Whenever the total is positive but the components are missing, I impute their values by using the industry average shares. I cannot impute missing values for the food and beverage industry (Nace 15) because no firm is reporting the subcategories.

Value Added: I drop the observation if it is negative.

Total labor cost (R216): when I take total employment from DADS I also replace total labor cost with the sum of gross salaries over all establishments from DADS (S_BRUT). I do so only when Total costs are non-missing, otherwise total costs would be heavily underestimated (would contain labor cost only).

Total costs: instead of purchases of goods (R210), purchases of raw materials (R212), and other purchases and charges (R214), I use their reported sum (total purchases and external charges, R771) whenever it is bigger.

Turnover (R310): if it is zero, it is set equal to the sum of turnover over all branches (S001) when this is positive. I also use turnover from branches if reported exports are larger than turnover but smaller than turnover from branches (only two cases in 2005).

Finally I drop the observations in the following cases:

- Value added is negative;
- Turnover comes entirely from branches classified in services;
- Turnover is lower than total exports. More precisely I allow for a 10% reporting error, hence I drop the observation if total exports are 10% bigger than turnover.

2.C Extra Results

Table 2.C.1: Outsourcing of Non-core Activities and Coordination Complexity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Exporter	0.090***								
	(0.017)								
NC		0.069***	0.084***	0.083***	0.080***	0.079***	0.079***	0.087***	0.090***
		(0.010)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.012)	(0.013)
Export			-0.012***	-0.012**	-0.010**	-0.010**	-0.011**	-0.011**	-0.010*
			(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Capital Intensity				0.007	0.011	0.011	0.012	0.013	0.006
				(0.011)	(0.011)	(0.011)	(0.012)	(0.013)	(0.013)
Skill Intensity					0.042***	0.041***	0.042***	0.050***	0.052***
					(0.010)	(0.010)	(0.010)	(0.012)	(0.012)
HQ Intensity						0.081**	0.081**	0.085**	0.087**
						(0.036)	(0.036)	(0.037)	(0.038)
Scale							0.009	0.000	-0.003
							(0.022)	(0.024)	(0.025)
Num imp. products								0.005	0.010
								(0.008)	(0.009)
Contract Intensity									0.011*
									(0.006)
Observations	237,085	185,964	185,964	184,879	176,260	176,038	176,038	152,245	146,251
Number of firms	39,595	31,287	31,287	31,147	30,246	30,234	30,234	26,168	25,417
R-sq O	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of outsourcing (Out_1b) over total costs. All variables are in logs apart from HQ Intensity. Data are for period 1996-2007. Firm-clustered standard errors in parentheses; (*, **. ***) indicate 10, 5, and 1 percent significance levels.

Table 2.C.2: Service Outsourcing and Country Re-entry

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Exporter	0.042***	0.034***	0.034***	0.034***	0.033***	0.034***	0.034***	0.034***	0.039***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.010)
NC		0.156***	0.156***	0.155***	0.148***	0.148***	0.145***	0.139***	0.135***
		(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.018)
NC-Reentry		-0.007	-0.007	-0.007	-0.006	-0.006	-0.006	-0.006	-0.001
		(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
NC#NC-Reentry		-0.004**	-0.004**	-0.004**	-0.004**	-0.004**	-0.003**	-0.003**	-0.004**
		(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Export			0.006	0.005	0.001	0.001	0.004	-0.000	0.003
			(0.011)	(0.011)	(0.012)	(0.012)	(0.015)	(0.014)	(0.014)
Capital Intensity				-0.000	-0.002	-0.002	-0.002	-0.002	-0.001
				(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Skill Intensity					0.001	0.001	0.001	0.001	0.002
					(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
HQ Intensity						-0.004	-0.004	-0.005	-0.003
						(0.005)	(0.005)	(0.005)	(0.005)
Scale							0.023	0.017	0.017
							(0.014)	(0.013)	(0.013)
Num imp. products								0.024**	0.029***
								(0.011)	(0.011)
Contract Intensity									0.003
									(0.008)
Observations	235,182	235,182	235,182	234,756	224,561	224,260	224,260	224,260	167,621
Number of firms	39,500	39,500	39,500	39,457	38,326	38,300	38,300	38,300	29,027
R-sq W	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Fixed effects	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr	ind#yr

Note: The dependent variable is the share of outsourcing over total costs. All variables are in logs apart from HQ Intensity. Data are for period 1996-2007. Firm-clustered standard errors in parentheses; (*, **. ***) indicate 10, 5, and 1 percent significance levels.

Chapter 3

Variety Growth, Welfare Gains and the Fall of the Iron Curtain

3.1 Introduction

The love-of-variety motif has been a building block of much of the literature on international trade since the seminal work of Krugman (1979). For the first time, Krugman showed how gains from trade can stem from the import of new varieties. Given the importance and the extensive application of this assumption, it is quite surprising that the work of Broda and Weinstein (2006, hereafter BW) is the first attempt to structurally estimate the impact of the change in variety on the whole economy. The authors extend Feenstra's (1994) aggregate price index to the case of several goods. Then, by using a similar empirical strategy, they estimate the import bias resulting from the omission of new varieties. From the bias they trace back the welfare gains due to variety growth, that is, how much consumers are willing to pay to access the larger set of varieties available at the end of the analyzed period. They show that the unmeasured growth in product variety has been an important source of gains from trade in the US over the 1972-2001 period. Indeed the import bias in the conventional price index is 28% or 1.2 percentage points per year and the welfare gains are equal to 2.6% of GDP.

In this study, I take a similar approach to estimate the welfare gains deriving from the import of new variety in the United Kingdom (UK). This is an intrinsically interesting question because no analogous estimates exist for the UK over the same time period. Moreover the paper offers new insights on the sources of variety growth in the case of the UK and Europe in general. BW find that the impact of variety growth is much stronger over the '70s and '80s but they do not provide an explanation for why this is the case and just refer to globalization in general. However the phenomenon of globalization has been constantly present over the past 30 years and, if anything, some of its driving forces have been even stronger in the last decade. Consider the surge of China and, in the case of the US, trade liberalizations like the North American Free Trade Agreement that came into effect in 1994. It is therefore not entirely clear why variety growth has such a stronger effect over the '70s and '80s.

This result may have been partially driven by the estimation procedure. In fact, BW estimate the welfare gains during the first part of the analyzed period using a different and more aggregated classification; intuitively goods defined at a more aggregated level are characterized by lower values of elasticities and this can bias the estimates of welfare gains upwards. The present analysis will prove that, although the bias is likely to be present, it is not possible to establish the direction of the bias a priori and even harder is to quantify its magnitude in practice. In this respect the case of the UK is a very interesting one: an important episode like the Fall of the Iron Curtain happened in the early '90s. Thus one could expect to find higher welfare gains during the later period compared to the '70s and '80s. By keeping a consistent estimation strategy throughout the period, I show that this is in fact the case, and that the direct contribution of former Soviet countries to the welfare gains is quantitatively important.

This paper therefore extends the work of BW by shedding light on the determinants of the welfare gains from variety growth, without simply ascribing these effects to the general phenomenon of globalization. Moreover I refine the strategy for handling the measurement error and for obtaining the weights in the weighted least squares estimator. Finally I provide an extensive robustness analysis, which assesses the effects of various assumptions untested in the BW's study: namely, the choice of the data and the classification, the definition of the set of goods, the weighting scheme to correct for heteroskedasticity and the form of the measurement error. Both the elasticity estimates and the final welfare gains will prove to be quite sensitive to the weighting scheme, in particular.

The paper is organized as follows. The next section outlines the theoretical framework that justifies the gains from new variety and provides a clear definition of variety; then in the subsequent section I present the data and aggregate statistics. Section 4 precisely describes the model behind the empirical strategy and provides a robustness analysis for the various assumptions. Section 5 presents the main results. Finally conclusions are outlined in Section 6.

3.2 Theoretical Background and Related Literature

Following BW's approach, the theoretical framework that underpins the welfare gains from variety growth is provided by the models of Spence (1976) and of Dixit and Stiglitz (1977). The success of this framework is due to the simple specification of how consumers value variety. The constant elasticity of substitution (CES) utility function proves indeed very useful in empirical studies. It is very tractable, it allows me to aggregate price changes across markets and, above all, it is fairly easy to estimate, having a simple demand structure.

Another important issue is the definition of variety. Intuition would suggest that a variety of a good should be defined as a specific product produced by a firm; for instance a particular brand of orange juice. This strategy would also allow me to adhere more closely to Krugman's (1980) model, whose structural assumptions are taken in this study. Unfortunately, this clashes with a sheer practical problem: obtaining bilateral firm-level export flows from all the countries in the world is infeasible. The final aim of this study is to calculate the welfare gains for the whole British economy, and so all possible imports have to be considered. Therefore the Armington's (1969) definition is adopted: goods traded internationally are considered differentiated on the basis of their country of origin. Thus a variety is simply a particular good produced by a particular country, e.g., Spanish oranges.

The drawbacks of this definition are evident. It is clearly possible that a good of the same brand is imported from two different countries: Adidas trainers imported from Vietnam and China for instance. Still, it ensures an immediate empirical application, with data coming from databases such as Comtrade or the OECD's International Trade by Commodity Statistics (ITCS). Moreover, the potential problems are mitigated by the empirical strategy adopted in the analysis. By estimating the elasticities of substitution at the good level, it is in fact possible to obtain a diversification in the degree of substitution of varieties. Whenever the elasticity of substitution for a particular good is high, consumers tend to be rather indifferent in choosing among different varieties, as in the case of commodities for instance. Consumers therefore do not differentiate in terms of the country of origin and the potential gains from variety growth are small, minimizing the problems related to the definition of variety. Low values of the elasticity of substitution on the other hand indicate that consumers care about the different varieties; hence, an increase in the number of supplying countries (i.e.: varieties) may constitute a source of welfare gains due to the love-of-variety motif introduced by the Krugman's model.

More recently, Feenstra (2009) and Feenstra and Weinstein (2010) have pointed out that the gains from new imported varieties can be partially offset by the welfare loss from fewer domestic varieties. On the other hand, they introduce another source of welfare gains: the reduction in firm markups due to import competition. In order to achieve this result, they assume translog preferences that allow for endogenous markups. The approach of BW may therefore overstate the gains from imported variety because it does not allow for domestic exit and because, conversely to the translog case, the CES system implies an infinite reservation price for varieties that are not consumed. Feenstra and Weinstein (2010) calculate the consumer gains for the US over the 1992-2005 period and find that only two-thirds of the gains come from variety growth, the remainder is a competition effect due to lower markups. Their approach would probably be the best way to proceed but unfortunately data constraints do not allow me to apply the same methodology to the UK. In any case, they find that the overall gains are comparable to the estimates of BW. The analysis of the present study therefore provides estimates for the UK that are sensible, at least in the aggregate.

Moreover, variety growth can still be underestimated in the present setting if the market concentration of the supplying countries decreases over time, due to the entry of new supplying firms that is not captured by industry data. Blonigen and Soderbery (2010) for instance show that welfare gains from new varieties in the US automobile sector are twice as large as standard estimates when detailed firm level data are used. In addition, Halpern et al. (2011) argue that losses due to import substitution may not be very large in practice: they find a relatively low elasticity of substitution between domestic and foreign intermediate goods, and the losses to domestic input suppliers caused by a trade liberalization are partially offset by increased demand for their products due to overall higher firm productivity. Finally, Feenstra and Weinstein (2010) adopt an estimation and weighting scheme strategy very similar to that of BW, hence the investigation of the robustness of such approach is still very relevant.

3.3 Data and Macro Evidence

In terms of data, I use three main sources: the UN Commodity Trade Statistics Database (Comtrade), the OECD's International Trade by Commodity Statistics (ITCS), and the Eurostat's Comext. All of them provide bilateral trade flows under several classifications. In this study three main classifications are used: the Standard International Trade Classification (SITC) Revision 2 and 3, the Harmonized System (HS) 1988, and the Combined Nomenclature (CN) 2005. In principle one would like to use data at the most disaggregated level in order to obtain a precise definition of goods; so the CN classification would be the obvious choice since it provides data with the most updated definition of goods and at the most disaggregated level (8 digit). However data availability problems and re-classification issues impose a trade-off. The problem of the Combined Nomenclature is that it is updated over time. The underlying classification is the HS, so the two nomenclatures coincide until the 6th digit, then the CN includes two other extra levels of disaggregation providing categories at the 8th digit. The HS classification was reviewed in 1996 and 2002 including new categories, and the CN has incorporated these changes. The definition of goods is therefore changing over time making it hard to track the same products over the whole time period. For this reason, the HS classification is taken as the benchmark and results obtained with the other classifications are compared against it.

However, this strategy has a number of limitations. First of all, restricting the analysis to the 6-digit level brings about a smaller variety growth. Indeed, the UK essentially imports all goods at that level of aggregation, which implies that the possibility of an increase in the number of varieties through the new good dimension is fairly small. The number of goods is constrained by the classification structure, because the total number of categories is bounded. New goods are initially classified in existing categories, which leads to an underestimation of variety growth. This limitation is immediately clear from Table 3.1. Compared to the case of the U.S. analyzed by BW, we can see that, using 6-digit level data, the number of goods remains almost constant over time. The table does show a clear increase in the number of countries supplying each individual good but the data basically hide any increase in the number of goods. Whilst the average number of supplying countries rises by 23% over the period to roughly 24 countries, the number of goods stays almost constant. In fact it even decreases to 4,815 from 4,886. This is a clear failure of the adopted classification and the first immediate consequence is an underestimation of variety growth since the channel of good expansion, which in BW's study counted for as much as half of the total increase in the number of varieties, is ruled out here.

Using the CN classification, the number of goods available is more comparable with the results for the US, which rely on the TSUSA/HTS systems. Table 3.2 shows that the number of categories is equal to 9,493 in 2006; this is still just 58% of total categories available for the HTS system but it is quite an improvement with respect to the HS classification. The total number of categories does not rise that much again, just by 4%; but there is a lot of churning, with many categories disappearing and new ones appearing. The total number of varieties increases by 31% over the period while the average number of supplying countries rises by 26%. So, compared to the HS classification, the number of both varieties and supplying countries grows

Table 3.1: Variety in UK Imports, Harmonized System, 1988-2006

Year	Goods	Number of HS categories	Average number of exporting countries	Total number of varieties	Share of total imports
1988	All	4,886	19.4	95,025	100%
	Common	4,714	19.8	93,445	94.6%
	Not in 2006	172	9.2	1,580	5.4%
2006	All	4,815	23.9	115,131	100%
	Common	4,714	24.1	113,590	94.2%
	Not in 1988	101	15.3	1,541	5.8%

Source: OECD - ITCS - HS 1988.

Table 3.2: Variety in UK Imports, Combined Nomenclature, 1988-2006

Year	Goods	Number of CN categories	Average number of exporting countries	Total number of varieties	Share of total imports
1988	All	9,129	14.8	135,425	100%
	Common	5,503	15.0	82,590	51.0%
	Not in 2006	3,626	14.6	52,835	49.0%
2006	All	9,493	18.7	177,339	100%
	Common	5,503	19.2	105,814	45.0%
	Not in 1988	3,990	17.9	71,525	55.0%

Source: Eurostat - CN 2009

slightly more. The main difference is that the number of average supplying countries per good is lower: 18.7 versus 23.9 in 2006. This fact is quite intuitive given the higher number of available categories. Overall, the data from the CN system are more comparable to those used by BW. Unfortunately, the problem is that this classification is much more unstable over time and, as will become clear from the next sections, this can be a source of concern because it negatively affects the definition and the number of goods that can be estimated.

Even though the figures in Table 3.1 and 3.2 are not directly comparable with those reported by BW for the US due to the different level of aggregation and different periods, it is still possible to compare the two economies from a broader perspective. It is well known that trade has a larger importance in the UK economy. The share of imports of goods in UK GDP is more than 80% bigger than the same value in the US: in 2001 it was 25.7% in the UK versus 14.1% in the US.[1] Still, imports have been growing faster in the US in the past few decades: over the period 1982-2006 the share of imports has risen by 84% in the US versus 33% in the UK. This growth is also reflected in the percentage increase in the number of supplying countries that has been higher in the US; in recent years the US has caught up with the UK in terms of average number of supplying countries.[2] In terms of supplying countries, the story is rather similar to the one in the US. Table 3.3 shows the surge of China, which moved from the twenty-third to the fifth place, the fairly good performance of India and the fall of Japan. A novel aspect is the evidence for the fall of the Iron Curtain; Russia moved from the twenty-fifth to the thirteenth place, Poland from the thirty-sixth to the twentieth and the Czech Republic from the forty-seventh to the twenty-eighth place despite the separation from the Slovak Republic.[3] Even more striking is the evidence from Figure 3.1. It displays the total growth in number of varieties versus the growth of varieties coming from the countries that formed the Soviet Bloc.[4] Total varieties have increased by a mere 20% (left axis) while varieties from ex-Soviet countries have almost tripled, increasing by a factor of 2.8. This growth has resulted in an analogous increase of the ratio of varieties imported from these countries to the total number of varieties (right axis): it has risen from 5% in 1988 to 11.5% in 2006. In terms of trade (Figure 3.2) the increase has been even sharper: total trade from ex-Soviet countries has increased by a factor of 13! And the actual growth is probably underestimated because the German Democratic Republic joined Germany in 1990 so is obviously excluded. Moreover the results shown in the figure are calculated with data from the HS classification and it might well be the case that, given the fixed classification, the actual growth in varieties has been even bigger.

[1] Data are from the World Bank World Development Indicators.

[2] BW report an average value of 15.8 in 2001 versus 18.7 for the UK in 2006. Considering the HS classification the number of supplying countries in the UK rises to 23.9 in 2006, but this is caused by the mechanical effect briefly outlined in the text.

[3] In Table 3.3, countries are ranked in terms of trade value. Little changes in terms of number of goods imported; Poland moved from the thirtyfourth to the twentieth place and Czech Republic from the thirty-second to the twenty-sixth; only Russia did not move, ranking forty-fourth both in 1988 and 2006.

[4] The former Soviet countries are: the USSR Republics (Armenia, Azerbaijan, Belarus, Estonia, Georgia, Kazakhstan, Kygyzstan, Latvia, Lithuania, Moldova, Russia, Tajikistan, Turkmenistan, Ukraine, Uzbekistan); the countries in the Comecon (Albania, Bulgaria, Czech Republic, Cuba, German Democratic Republic, Hungary, Mongolia, Poland, Romania, Social Republic of Vietnam) and the associate countries (Yugoslavia that split into: Bosnia, Croatia, Kosovo, Macedonia, Serbia & Montenegro, Slovenia).

Table 3.3: Ranking of UK Supplying Countries

Country	1988	1993	2006
Germany	1	1	1
United States	2	2	2
France	3	3	3
Netherlands	4	4	4
Japan	5	5	11
Italy	6	6	8
Belgium	7	7	7
South Africa	8	11	17
Ireland	9	8	10
Sweden	10	10	12
Norway	11	9	6
Spain	12	12	9
Denmark	13	15	14
Switzerland	14	14	19
Canada	15	17	15
Finland	16	16	24
Hong Kong	17	20	27
Chinese Taipei	18	18	25
Korea	19	25	22
Portugal	20	23	23
Austria	21	26	26
Brazil	22	28	33
China	23	13	5
Australia	24	27	30
Former USSR	25	31*	13*
India	26	24	21
Singapore	27	21	18
New Zealand	28	34	53
Saudi Arabia	29	22	47
Israel	30	33	39
Poland	36	35	20
Former Czechoslovakia	47	42**	28**

Supplying countries are ordered in terms of trade value; top 30 countries in 1988 included, plus Poland and Czech Republic. Data from OECD-ITCS HS 1988. *: Russian Federation; **: Czech Republic

It is clear that, although the HS classification allows me to track the goods over the entire period, this choice comes at a cost. The growth in variety is probably underestimated and this can jeopardize the outcome of the analysis. In the next sections I will therefore show results obtained with different classifications to get a sense of the sensitivity to this choice. Moreover, from this first analysis, it seems clear that the fall of the Iron Curtain might well have had a significant impact on variety growth. Or, even more ambitiously, it could be claimed that this event has been one of the prime causes of variety growth and the consequent welfare gains in the UK over the past twenty years.

3.4 Empirical Strategy

The empirical strategy of this paper is based on the work of BW that, in turn, draws on the work of Feenstra (1994). The main merit of BW's analysis is to generalize Feenstra's methodology to calculate an aggregate import price index by combining all imported goods into a composite import good. In this study, I apply the same estimation method and perform a robustness analysis to compare different versions of the estimation model in terms of the weighting scheme and of the strategy to handle the measurement error.

3.4.1 Theory

BW extend the Feenstra Price Index, which allows for taste and variety changes for a single good, to the case of several CES aggregate goods. They adopt a three-level CES utility function, similar to Helpman and Krugman (1985, Ch. 6), which allows them to estimate the impact of variety growth on prices and then welfare. In order to understand the main contributions of the two papers, it is necessary to define the main quantities that pin down the exact price index. The utility function is separable into a domestic good, D_t, and a composite imported good, M_t. The composite imported good gives the overall utility at time t generated by the consumption of all imported goods. A particular imported good, g, is in turn composed of all the varieties of that good consumed at time t. The lowest level of the utility function, which coincides with the utility derived by the consumption of a single good, and the corresponding minimum unit-cost function are defined as follows:

$$M_{gt} = \left(\sum_{c \in C} d_{gct}^{1/\sigma_g} (m_{gct})^{(\sigma_g - 1)/\sigma_g} \right)^{\sigma_g/(\sigma_g - 1)} \tag{3.1}$$

$$\Phi_{gt}^{M}(I_{gt}, \mathbf{d}_{gt}) = \left(\sum_{c \in I_{gt}} d_{gct} (p_{gct})^{1-\sigma_g} \right)^{1/(1-\sigma_g)} \tag{3.2}$$

where m_{gct} is the particular variety of good g imported from country c at time t; σ_g (>1) is the elasticity of substitution among varieties of good g; d_{gct} is a taste or quality parameter; p_{gct} is the price (or better unit value) of variety c of good g in period t. C is the set of all countries and $I_{gt} \subset C$ is the subset of all varieties of good g consumed in period t.

The second level of the utility function, which aggregates over all goods and pins down the composite imported good, M_t, is similarly defined, together with the unit-cost function, as follows:

$$M_t = \left(\sum_{g \in G} M_{gt}^{(\gamma-1)/\gamma} \right)^{\gamma/(\gamma-1)} \tag{3.3}$$

$$\Phi_t^M = \left(\sum_{g \in G} (\Phi_{gt}^M (I_{gt}, \mathbf{d}_{gt}))^{1-\gamma} \right)^{1/1-\gamma} \tag{3.4}$$

where γ (>1) denotes the elasticity of substitution among imported goods and $\mathbf{d}_{gt}$ is the vector of quality parameters for each country. Finally the upper level of the utility function that aggregates the domestic good and the composite imported good is given by:

$$U_t = \left(D_t^{(\kappa-1)/\kappa} + M_t^{(\kappa-1)/\kappa} \right)^{\kappa/(\kappa-1)} \tag{3.5}$$

where κ (>1) is the elasticity of substitution between the two goods.

We are now ready to state the two main contributions of BW and Feenstra's works in the following two propositions.

Proposition 1. For g $\in G$, if d_{gct}=d_{gct-1} for c $\in I_g = (I_{gt} \cap I_{gt-1})$, $I_g \neq \varnothing$, then the exact price index for good g with change in varieties is given by:

$$\pi_g = \frac{\Phi_{gt}^M (I_{gt}, \mathbf{d}_{gt})}{\Phi_{gt-1}^M (I_{gt-1}, \mathbf{d}_{gt})} = P_g^M \left(\frac{\lambda_{gt}}{\lambda_{gt-1}} \right)^{1/(\sigma_g - 1)} \tag{3.6}$$

where $\lambda_{gt} = \frac{\sum_{c \in I_g} p_{gct} x_{gct}}{\sum_{c \in I_{gt}} p_{gct} x_{gct}}$, $\lambda_{gt-1} = \frac{\sum_{c \in I_g} p_{gct-1} x_{gct-1}}{\sum_{c \in I_{gt-1}} p_{gct-1} x_{gct-1}}$ and P_g^M is the conventional price index for good g over a constant set of varieties.

Proposition 2. If d_{gct}=d_{gct-1} for c $\in I_g \neq \varnothing \ \forall g \in G$, then the exact **aggregate** import price index with variety change is given by:

$$\Pi_g = \frac{\Phi_t^M (I_t, \mathbf{d}_t)}{\Phi_{t-1}^M (I_{\mathbf{t-1}}, \mathbf{d}_t)} = CIPI \prod_{g \in G} \left(\frac{\lambda_{gt}}{\lambda_{gt-1}} \right)^{\omega_{gt}/(\sigma_g - 1)} \tag{3.7}$$

where $CIPI$ is the aggregate conventional import price index, $\omega_{gt} = f(p_{gct}, x_{gct})$ are log-change ideal weights.

Although rather condensed, it is quite straightforward to understand the main results of the two propositions that pin down the welfare gains stemming from variety growth. In summary, the main goal of the analysis is to compute the value of the lambda ratio for each good, which is the deviation of the exact price index from the conventional price index. The lambda ratio defines the importance of new varieties: the higher the expenditure share on new varieties, the lower is λ_{gt}, and the smaller is π_g with respect to the the conventional price index P_g^M.

The lambda ratio also depends on σ_g, which is the estimated elasticity of substitution for the particular good g. It is important to notice that when the elasticity is big the lambda ratio tends to one so the difference between the two price indices is small. This implies that when new varieties are close substitutes to the existing ones, the Feenstra Price Index does not differ much from the conventional price index and the gains from variety growth are small because consumers do not care much about the new varieties. So, in this framework, the measure of variety growth is not simply given by the number of varieties but it takes into account taste or quality differences that affect the share of expenditures among the different varieties. This corrects the so-called "quality bias". Moreover, allowing for good-specific values of the elasticity of substitution, it is also possible to correct the "symmetry bias" among the available goods.

Then the difference between the exact and the conventional aggregate import price index is simply calculated as a geometric weighted average of the lambda ratios, this term is referred to as import bias and for clarity the formula is reported below:

$$IB = \prod_{g \in G} \left(\frac{\lambda_{gt}}{\lambda_{gt-1}} \right)^{\omega_{gt}/(\sigma_g - 1)} \tag{3.8}$$

The weights are ideal log-change weights, which are a function of prices and quantities for all the varieties of a particular good. Finally, the welfare gains due to variety growth are obtained by raising the import bias to the ideal import share over the considered period; this share represents the fraction of imported goods in total GDP.[5] It is worth noticing that the import bias is defined over the period taken into consideration; therefore in Proposition 1 and 2 one should read the starting and final year of the period, instead of *t* and *t-1*.

3.4.2 Estimation Method

The entire procedure for obtaining an estimate of the welfare gains due to variety growth can be summarized by the following steps:

1. Define the set of goods G;

2. Obtain estimates of the good-specific elasticity of substitution, σ_g;

3. Calculate the λ_{gt} ratios, which capture the role of new varieties for every good g;

4. By combining estimates of σ_g with the measures of variety growth for each good, obtain an estimate of how much the exact price index for good g changes as a result of the change in varieties (the lambda ratio);

5. Apply the ideal log weights (ω_{gt}) to the price movements of each good in order to obtain an estimate of the bias on the exact aggregate price index (the import bias, see equation 3.8);

[5]All the log-change ideal weights are defined in the appendix in equations (3.A.2) and (3.A.3).

6. Calculate the welfare gain or loss from these price movements using the ideal import share in the period (equation 3.A.3);

7. Bootstrap the entire procedure to obtain an estimate of the standard error of the various quantities.

The first two steps are the most involved ones; the definition of the set of goods is not as obvious as one might think. It is indeed affected by the estimation methodology of σ_g, which, overall, is the core part of the entire procedure. The main steps and assumptions that underline the estimation of the elasticities of substitution are shown in the Appendix. At this stage, it suffices to understand the final equation that takes the following form:

$$(\Delta^k \ln p_{gct})^2 = \theta_1(\Delta^k \ln s_{gct})^2 + \theta_2(\Delta^k \ln p_{gct}\Delta^k \ln s_{gct}) + u_{gct} \tag{3.9}$$

where $\Delta^k x_{gct} = \Delta x_{gct} - \Delta x_{gkt} = (x_{gct} - x_{gct-1}) - (x_{gkt} - x_{gkt-1})$, s_{gct} is the cost share of variety c in total expenditures of good g (defined in equation 3.A.5), and $\sigma_g = f(\theta_1, \theta_2)$ as shown in Proposition 3 in the Appendix. This equation is obtained by assuming a particular form for the supply curve and taking first differences of prices (unit values) and shares twice: first with respect to time and then to a reference country k. Unfortunately it still not possible to consistently estimate equation (3.9) because prices and expenditure shares are correlated with the error term. Nevertheless, it is possible to obtain a consistent estimator for the thetas and hence for the elasticity of substitution by averaging (3.9) over time. The estimation is still possible because σ_g and the supply elasticity are assumed to be constant over the varieties of the same good; the former due to the CES demand structure, the latter for the particular form of the supply curve, whose elasticity is assumed to be equal across all supplying countries.[6]

Hence, taking the sample means of the variables, equation (3.9) can be rewritten as:

$$\overline{(\Delta^k \ln p_{gct})^2} = \theta_1\overline{(\Delta^k \ln s_{gct})^2} + \theta_2\overline{(\Delta^k \ln p_{gct}\Delta^k \ln s_{gct})} + \overline{u_{gct}} \tag{3.10}$$

From the assumption that underlines the identification strategy, $E\left[\overline{u_{gct}}\right] = 0$.[7] This implies that the expectation of the error term in (3.10) converges to zero and the equation can be consistently estimated. Let $\widehat{\theta}_1$ and $\widehat{\theta}_2$ denote the estimates of θ_1 and θ_2 obtained by running weighted least squares (WLS) on (3.10), it turns out that these are equivalent to the Hansen's (1982) GMM estimator defined as follows:

$$\widehat{\beta}_{GMM} = \arg\min_{\theta\in\Theta} \overline{u_{gct}}(\beta)' W \overline{u_{gct}}(\beta) \tag{3.11}$$

where $\beta = \binom{\sigma_g}{\rho_g}$ as defined in equation (3.A.13), and W is a positive definite weighting matrix. In order to get a consistent estimate for σ_g, $\widehat{\theta}_1$ and $\widehat{\theta}_2$ are first obtained by running WLS on (3.10); then σ_g is computed using Proposition 3. Whenever an unfeasible value for σ_g is obtained

[6] The supply elasticity is defined in equation (3.A.7).

[7] Supply and demand error terms are assumed to be independent, see equation (3.A.10). A further condition to get identification requires to have some difference in the relative variances of the two error terms. See Feenstra (1994).

(<1), a constrained numerical minimization of equation (3.11) is performed using the Nelder and Mead's (1965) simplex algorithm.[8]

3.4.3 The Impact of the Good Definition

The definition of the set of goods would seem immediate. It is indeed intuitive to define goods with the most disaggregated available categories, in the present analysis at the 6 or 8-digit level. As shown in Table 3.1 and 3.2, there are 4,714 HS or 5,503 CN categories that are available over the 1988-2006 period, but unfortunately the elasticity cannot be estimated for all categories. The estimation strategy imposes to take first differences twice (with respect to time and to a reference country) and a minimum number of observations for each category is required to estimate the elasticity. Two years are lost by taking first differences with respect to time and by calculating the expenditure shares, the latter because the expenditure shares are defined on the set of common varieties between year t and year *t-1* so the first data point is lost in this calculation (see equation 3.A.5). Moreover the first difference with respect to the reference country k requires each good to have at least one country (i.e. variety) always be present in the data set, without any missing year, as stressed by Feenstra (1994). Finally, since there are two independent variables in equation (3.10), three countries for each good are at least needed in order to get identification. One of these has to be a supplier of the good in every year and the other two in at least two consecutive years.[9] This requirement affects the definition of the set of goods because it is not satisfied for all 6 or 8-digit categories. Whenever this is the case, a good is defined at a more aggregated level.

A very similar strategy is used by BW, with quite severe consequences on the number of final estimated goods; for the period 1990-2001 the number of final goods for which they are able to calculate the lambda ratios drops to 921 from the 14,549 available categories. In this study, the number of goods turns out to be equal to 2,466 for the HS classification and 1,079 for the CN system. Hence, even if the CN starts with many more categories, the final number of estimated goods is actually smaller due to the evolution of the classification. A difference with respect to BW's work is that I estimate the elasticities of substitution for the same goods for which the lambda ratios are calculated. Conversely to BW who use a weighted average of the more disaggregated elasticities whenever they have a good defined at a more aggregated level. The advantage of my approach is that I will use the exact same underlying data for the estimation of elasticities and the calculation of the lambda ratios, and I will be able to test more precisely the impact of the aggregation level at which the elasticities are estimated on the final welfare gains.

[8]The variables are constrained as follows: $1 < \sigma_g \leq 135.5$ and $0 \leq \rho_g < 1$. Once a solution is obtained the non-linear condition $\rho_g \leq (\sigma_g - 1)/\sigma_g$ is checked. If the condition is not satisfied, it implies that a bigger value of σ_g would be needed. In the rare event that this happens (on average less than 0.1% of all estimates), σ_g is assumed to be equal to 140. All results are very robust to these assumptions, the choice of a max value of sigma equal to 140 has no impact on the final result.

[9]The required number of countries is actually four, as it will become clear in the next section. In order to minimize the impact of measurement error, another variable is in fact added to the right hand side of equation (3.10).

Despite estimating elasticities at the most disaggregated level, the level of aggregation might still have an impact on BW's work because they use two different datasets. For the earliest period (1972-1988), they use the TSUSA system that has categories of goods defined at a more aggregated level (7 digits) with respect to HTS (10 digits), which they employ for the later period (1990-2001). The impact of this choice remains untested in their study but it may well have an impact on the final results. The elasticities of more aggregated goods are likely to be smaller, hence the calculated welfare gains could be biased upwards and this might explain why the find much larger welfare gains in the earlier period. I will therefore explicitly test the sensitivity of the welfare gains estimates to the level of disaggregation at which both elasticities and goods are defined.

Given the non-linearity of the problem, the standard errors on the estimates of the lambda ratio, the import bias and the welfare gains are obtained by bootstrapping each estimate of σ_g 50 times and by recalculating the various quantities for each set of parameters. As robustness check, standard errors on the σ_g are also calculated using the delta method for the sub-sample of goods that can be estimated analytically.

3.4.4 Measurement Error and Robustness to the Weighting Scheme

Unfortunately a direct measure of prices is not available, so p_{gct} is calculated as unit value. This procedure implies that prices are surely measured with some error. In order to mitigate this problem, Feenstra (1994) suggests to add another variable to equation (3.10); he simply adds a constant that will capture the variance of the measurement error. BW refine his method by making some extra assumptions on the form of the error, which also affects the strategy for the form of the weighting matrix in the WLS estimator.

Let p_{gcti} be the price of a particular product of variety c of good g; so that the trade value $p_{gct}x_{gct} = \sum_i p_{gcti}$ because the quantity of each product, x_{gcti}, always equals one (i.e.: in case of more items of the same product, the same price is added several times). They assume that product prices are measured with an i.i.d. error such that $p_{gcti} = \widetilde{p}_{gcti}\zeta_{gcti}$ where $\widetilde{p}_{gcti}$ is the true price and p_{gcti} is the measured price. In this case the error has mean zero and:

$$var(\ln \zeta_{gcti}) = \sigma^2 \tag{3.12}$$

$$cov\left(\ln \zeta_{gcti}, \ln \zeta_{gc^*sj}\right) = 0 \;\; \forall\, c \neq c^*,\; t \neq s,\; i \neq j \tag{3.13}$$

By assuming that the log of the geometric mean price of a variety is approximately equal to the log of the arithmetic mean, it is possible to compute the variance of $\ln p_{gct}$ as follows:

$$\sigma^2_{\ln p_{gct}} \equiv var\left[\ln\left(\frac{\sum_i p_{gcti}}{x_{gct}}\right)\right] \approx var\left[\ln\left(\left(\prod_i p_{gcti}\right)^{1/x_{gct}}\right)\right] =$$

$$= \frac{1}{x^2_{gct}} var\left(\sum_i (\ln \widetilde{p}_{gcti} + \ln \zeta_{gcti})\right) = \frac{1}{x^2_{gct}} x_{gct}\sigma^2 = \frac{1}{x_{gct}}\sigma^2 \qquad (3.14)$$

Thanks to the assumptions on structure of the error (zero mean and condition 3.13):

$$E(\Delta^k \ln p_{gct})^2 = E\left((\ln p_{gct} - \ln p_{gct-1}) - (\ln p_{gkt} - \ln p_{gkt-1})\right)^2 =$$

$$= \delta^2_{gc_kt} + \sigma^2\left(\frac{1}{x_{gct}} + \frac{1}{x_{gct-1}} + \frac{1}{x_{gkt}} + \frac{1}{x_{gkt-1}}\right) \qquad (3.15)$$

where $\delta^2_{gc_kt}$ is the variance of the true price differences over time and with respect to variety k. Averaging this across all periods:

$$E\frac{1}{T}\sum_t(\Delta^k \ln p_{gct})^2 = \frac{1}{T}\sum_t \delta^2_{gc_kt} + \sigma^2\frac{1}{T}\sum_t\left(\frac{1}{x_{gct}} + \frac{1}{x_{gct-1}} + \frac{1}{x_{gkt}} + \frac{1}{x_{gkt-1}}\right) \qquad (3.16)$$

This implies that the equation (3.10) should be modified by adding the following error adjustment term to the right hand side:

$$err\_adj\_2 = \widehat{\theta}_3\frac{1}{T}\sum_t\left(\frac{1}{x_{gct}} + \frac{1}{x_{gct-1}} + \frac{1}{x_{gkt}} + \frac{1}{x_{gkt-1}}\right) \qquad (3.17)$$

where $\widehat{\theta}_3 = \widehat{\sigma^2}$ is a parameter to be estimated. This equation slightly modifies BW's approach since they add an analogous term, given by:

$$err\_adj\_1 = \widehat{\theta}_3\frac{1}{T}\sum_t\left(\frac{1}{x_{gct}} + \frac{1}{x_{gct-1}}\right) \qquad (3.18)$$

which does not take into account the first difference with respect to the reference country k. Since an error is likely to be present also in the measure of the prices of variety k, the choice of the reference country might affect the final results, and the estimates are likely to be more robust to measurement error if the modified term err_adj_2 is added instead of the term err_adj_1.

BW use a similar line of reasoning to justify their strategy for the weighting matrix of the WLS estimator. Heteroskedasticity is very likely to be present because if prices are measured with error, so are their sample variances. Hence, they try to correct for this heteroskedasticity by assuming that the sample variances are inversely related to the quantity of goods used in order to calculate unit values and to the number of periods. They claim that the variance of the left-hand side of equation (3.10) is proportional to:

$$\frac{1}{T^3}\left(\frac{1}{x_{gct}}+\frac{1}{x_{gct-1}}\right) \tag{3.19}$$

Analogously to the analysis performed for the measurement error, the variance can be alternatively defined in order to account for the first differences with respect to the reference country k as well; the claim is that it takes the following form:

$$\frac{1}{T^3}\left(\frac{1}{x_{gct}}+\frac{1}{x_{gct-1}}+\frac{1}{x_{gkt}}+\frac{1}{x_{gkt-1}}\right) \tag{3.20}$$

and hence the weights are given by:

$$weight\_2 = T^{3/2}\left(\frac{1}{x_{gct}}+\frac{1}{x_{gct-1}}+\frac{1}{x_{gkt}}+\frac{1}{x_{gkt-1}}\right)^{-1/2} \tag{3.21}$$

The weights corresponding to the BW's form of the variance (equation 3.19) are similarly defined and are denoted as $weight\_1$.

I perform an extensive robustness analysis to assess how the various assumptions on the measurement error and on the weighting scheme affect the final results. Four different specifications of the model are defined on the basis of the error adjustment term and the weighting scheme.[10] In order to assess the presence of heteroskedasticity, three different tests are performed: the White's (1980) general test that does not impose any structure on the heteroskedasticity and two different versions of the Breusch and Pagan's (1979) test, which instead requires a specific alternative hypothesis on the nature of the heteroskedasticity. The first version is run by assuming that the variance is a function of the quantities defined in equation (3.19) or (3.20) depending on the model.[11] In the second version the variance is assumed to be simply proportional to the first independent variable in equation (3.10), that, in turn, is proportional to the expenditure shares.

Table 3.4 contains the results for the various specifications. The first thing to notice is that the number of goods is limited to the sub-sample for which the elasticity can be estimated analytically. The hypothesis of homoskedasticity can be immediately rejected; the general White test reveals some form of heteroskedasticity in all models, with significant estimates for more than 70% of the goods. For the analysis of the Breush-Pagan test, it makes more sense to look at the two unweighted specifications only (Model 0 and 2); a form of the variance, as assumed in equation (3.19) or (3.20), does not seem a completely implausible assumption, with 34% of significant estimates.[12] The two forms of variance perform in a very similar way. However, a variance proportional to expenditure shares seem to perform even better, with 40% of significant estimates. Among the two weighted models, results do not give a strong preference, even

[10] Model 0: *err_adj_1* and no weights; Model 1: *err_adj_1* and *weight_1*; Model 2: *err_adj_2* and no weights; Model 3: *err_adj_2* and *weight_2*.

[11] In models 0 and 2 the test is run against the first form of variance while in models 1 and 3 against the second one.

[12] With data from the SITC Rev. 3 classification, the number of signficant estimates increases to 40%. The results for the weighted models are also consistent: very few significant estimates are found for the Breush-Pagan test that assumes the same form of variance whilst a lot for the other version.

Table 3.4: Heteroskedasticity Tests

Statistic	Model (0)	Model (1)	Model (2)	Model (3)
Num. Of Goods	1939	2022	1919	2021
White Test	961***	1175***	942***	1003***
	263**	330**	267**	316**
	132*	142*	138*	168*
Total Significant	*69.93%*	*81.45%*	*70.19%*	*73.58%*
Breush-Pagan Test (Variance)	515***	25***	507***	25***
	92**	83**	94**	86**
	56*	172*	56*	187*
Total Significant	*34.19%*	*13.85%*	*34.24%*	*14.75%*
Breush-Pagan Test (Exp. Shares)	485***	1360***	489***	1281***
	176**	152**	179**	190**
	140*	74*	142*	79*
Total Significant	*41.31%*	*78.44%*	*42.21%*	*76.69%*

Data: OECD - ITCS - HS 1988; Period: 1988-2006. ***, **, * indicate significance at 1%, 5% and 10% level, respectively. Model 0: err_adj_1 and no weights. Model 1: err_adj_1 and weight_1. Model 2: err_adj_2 and no weights. Model 3: err_adj_2 and weight_2

though model 3, which is characterized by the modified form of variance (equation 3.20), performs slightly better in all tests. The choice of this specification is also suggested by the more conservative results that it generates, as it will be clear in the next section.

In principle the weighting strategy should not influence the estimator that should remain consistent regardless of the weights used, as long as the weights are uncorrelated with the error term. If the right structure is chosen, the estimator is more efficient but, ultimately, this does not affect the current analysis directly, since the final standard errors are obtained by bootstrapping the entire procedure. It turns out that the weights affect the value of the estimates significantly; the impact on the estimates of the elasticities is analyzed here, while the impact on the final welfare gains is postponed to the next section.

Table 3.6 reports the sample statistics for the elasticity estimates over the 1988-2006 period using the Harmonized System, the benchmark classification for this study. Comparing the unweighted models to the weighted ones (model 0 vs 1 and 2 vs 3), it is immediately clear that the weighting scheme affects the elasticity estimates significantly, as also Figure 3.3 shows. The hypothesis that the mean and the median are the same is statistically rejected, the unweighted models persistently display higher values of sigmas compared to the weighted ones. This result is extremely robust to the data and the classification used.[13] The difference between the two weighted models (model 1 vs 3) is not as sharp; the values of the mean and the median are slightly higher using the new proposed weights, i.e. model 3, but it is not possible to statistically reject the hypothesis that the values are the same.[14] Still, this result persists when different data

[13]See the analogous estimates in Table 3.5, 3.B.1 and 3.B.2.

[14]The tests that the mean and the median are higher under model 3 have a p-value of 30% and 31%, respectively.

or classifications are used. Considering the elasticity estimates of Table 3.B.1, the hypothesis that the median is the same tested against the alternative that it is higher under model 3 can be rejected at the 5% significant level. The same hypothesis can be rejected for the mean as well, but this result is sensitive to the exclusion of the two outliers. In the case of the estimates obtained with the CN classification (Table 3.5) the same hypothesis can be rejected both for the mean and the median at the 1% significant level. Hence the model characterized by the new proposed weights generates estimates of the elasticities that are generally higher, so in between the OLS estimates and the estimates obtained with the BW's original weights.

Several reasons can be advanced to explain the sensitivity of the elasticity estimates to the weighting scheme: a misspecification in the original model, a wrong definition of the weights or small samples can all be plausible explanations. The last one is quite likely, the error term in equation (3.10) in fact vanishes only if T goes to infinity. Since the sample contains at most 17 years, it might well be the case that the disturbances do not disappear.[15] Hence both the regressors and the weights are likely to be correlated with the error term and it is hard to assess whether the introduction of the weights is beneficial or introduces further distortions into the model. For this reason it is difficult to find evidence in favour of one model a priori, and, even though the weighting schemes are reasonable, dismissing immediately the OLS estimates could be quite dangerous.[16] I therefore carry out an extensive robustness analysis and obtain the results for all the models, not only for the benchmark model characterized by the new proposed weights. This allows me to evaluate the sensitivity of the welfare gains estimates to the choice of the weighting scheme, the good definition and the classification used.

Besides the econometric assumptions like the weighting scheme, the estimates of the elasticities can also be affected by the methodology itself. Equation (3.9) is obtained by taking first differences over time but from a welfare point of view it might be more sensible to consider price and expenditure share differences over longer spans of time. Frictions are likely to be present in the short run and, after all, a welfare analysis consists in comparing different long run equilibria. A first attempt to obtain estimates of elasticities that are better suited to analyze long run changes is to take differences over longer time spans, for instance 5 years instead of 1 year. The results obtained after modifying equation (3.9) accordingly and re-estimating the elasticities are quite interesting. Figure 3.4 shows how the distribution of the estimates change when 5-year differences are taken for model 1 (sigma_1_LR) and 3 (sigma_3_LR), and contrasts them with the estimates previously obtained. The median elasticity is significantly higher, the overall distribution is more spread and the right tail is thicker. Overall the values of elasticities are higher across the board and this can well have an impact on the welfare gains as the next section will show.

[15]See Soderbery (2010) for an analysis of the sensitivity of the elasticity estimates to small sample biases.

[16]As Angrist and Pischke (2008b) put it: "Any efficiency gain from weighting is likely to be modest, and incorrect or poorly estimated weights can do more harm than good" (p. 96).

3.5 Results

3.5.1 Welfare Gains in the UK and Robustness to the Good Definition

Since heteroskedasticity is very likely to be present in the data but the real weights are unknown, it is useful to obtain the welfare gains under several specifications in order to get a sense of the sensitivity of the estimates. The first set of results refers to the period 1988-2006 and it is generated using the CN classification. This classification varies over time and allows for the expansion in the number of goods, so it does not underestimate variety growth as in the case of the HS classification, whose categories are stable over time. The potential problem with this classification is related to the definition of the goods: the categories are updated over time and many goods cannot be identified at the 8-digit level. Hence the data have to be aggregated up and the goods must be defined at a more aggregated level, losing part of the appeal of a more detailed classification. The same set of models defined in the previous section is used; this exercise extends the earlier robustness analysis to get also a sense of how the final results vary depending on the assumptions about the measurement error and the form of the variance. Standard errors and confidence intervals are obtained by bootstrapping the entire procedure 50 times.

Table 3.5 summarizes the results for the main variables of interest. As one could have expected from the elasticity estimates, the results are not very robust to the weighting scheme; the reasons outlined in the previous section apply. Model 1, characterized by the weighting scheme proposed by BW, gives the lowest elasticity estimates whilst the unweighted models the highest. The proposed modified weighting scheme (model 3) lies in the middle, proving to be a sensible and possibly even more robust alternative. The welfare gains under model 3 amount to 3.80% over the whole period, not statistically different from the estimate of 3.62% obtained with model 1; the two unweighted models display smaller gains due to the sensibly higher elasticity estimates. The main reason why the gains are much bigger compared to the US is to be sought into the higher openness to trade of the UK; in fact the ideal import share is much bigger, namely 22.5% for the period 1988-2005 versus 10.3% in the US for 1990-2001.

Table 3.6 shows the analogous results obtained with the HS classification for the period 1988-2006. It is easy to note that the welfare gains are lower across the board. This is likely to be the consequence of the fixed definition of the categories in the HS system, which prevents the number of goods to increase over time. Hence the welfare gains due to variety growth are likely to be underestimated. The results are still quite sensitive to the weighting scheme. The two weighted models now give welfare gains around 2% over the period; while the welfare gains under the two unweighted models are sensibly lower, around 1.1%. Again, the welfare gains obtained with the modified weighting scheme are slightly higher than the gains estimated with model 1, despite the higher values of the mean and the median of the elasticity estimates. As pointed out in the previous section, the hypothesis that the mean and the median are the same cannot be statistically rejected and the values of the 95th percentile suggest that the higher welfare gains under model 3 are explained by lower elasticity estimates on the right tail of the distribution. This result contrasts with the findings obtained with the SITC classification (Table

Table 3.5: The Impact of Variety Growth in the UK, 1988-2006, Combined Nomenclature

Quantity	Statistic	Model (0)	Model (1)	Model (2)	Model (3)
Num. Of Goods		1079	1079	1079	1079
2 digit		2	2	2	2
3 digit		111	111	111	111
4 digit		133	133	133	133
5 digit		64	64	64	64
6 digit		65	65	65	65
7 digit		20	20	20	20
8 digit		684	684	684	684
Sigma	Mean	13.80	8.72	14.01	11.89
	St. err.	2.14	0.79	1.02	0.96
	Percentile 5	1.87	1.46	1.91	1.56
	Median	3.61	2.77	3.75	3.01
	St. err.	0.07	0.06	0.08	0.06
	Percentile 95	135.49	20.56	135.50	135.50
Lambda Ratio	Percentile 5	0.604	0.464	0.614	0.526
	Median	0.987	0.979	0.988	0.982
	Percentile 95	1.309	1.552	1.265	1.461
Bias	Estimate	0.878	0.861	0.879	0.854
	Conf. Int.	[0.850, 0.901]	[0.759, 0.902]	[0.862, 0.893]	[0.757, 0.897]
Welfare Gains	Estimate	3.14%	3.62%	3.12%	3.80%
	Conf. Int.	[2.52%, 3.92%]	[2.47%, 6.77%]	[2.73%, 3.58%]	[2.61%, 6.81%]
Soviet Countries	2006	11.67%	11.31%	11.69%	13.87%
	Mean	9.52%	9.63%	9.54%	11.49%

Data: Eurostat - Comext - CN 2009. See notes in Table 4.

Table 3.6: The Impact of Variety Growth in the UK, 1988-2006, Harmonized System

Quantity	Statistic	Model (0)	Model (1)	Model (2)	Model (3)
Num. Of Goods		2466	2466	2466	2466
3 digit		65	65	65	65
4 digit		138	138	138	138
5 digit		117	117	117	117
6 digit		2146	2146	2146	2146
Sigma	Mean	15.75*	8.28	16.60*	8.78
	St. err.	0.87	0.51	1.13	0.79
	Percentile 5	1.74	1.40	1.73	1.50
	Median	3.45	2.70	3.47	2.72
	St. err.	0.05	0.04	0.06	0.03
	Percentile 95	135.50	19.62	135.50	18.99
Lambda Ratio	Percentile 5	0.632	0.488	0.634	0.536
	Median	0.979	0.965	0.979	0.967
	Percentile 95	1.094	1.147	1.092	1.131
Bias	Estimate	0.956	0.926	0.956	0.917
	Conf. Int.	[0.943, 0.963]	[0.885, 0.949]	[0.943, 0.963]	[0.843, 0.946]
Welfare Gains	Estimate	1.08%	1.85%	1.07%	2.08%
	Conf. Int.	[0.91%, 1.40%]	[1.24%, 2.94%]	[0.90%, 1.40%]	[1.32%, 4.15%]
Soviet Countries	2006	11.11%	8.62%	11.08%	10.37%
	Mean	8.36%	6.44%	8.34%	7.80%

Data: OECD - ITCS - HS 1988. *: one outlier is excluded, value of 8,945 in (0) and 36,115 in (2). See notes in Table 4.

3.B.2) where the results are more intuitive: the elasticity estimates are higher under model 3 and consequently the final welfare gains are lower. The reason seems to be again related to the elasticity estimates on the right tail: this time they are larger, as for the mean and the median. These results highlight how the final welfare gains depend on the shape of the entire distribution of the elasticity estimates, and not just the mean or the median.

As the variety growth is likely to be underestimated under the HS, it is quite safe to take 2% as a benchmark for the welfare gains over the period; this is larger than the smallest estimate but it is still far from the upper estimate obtained with the CN classification, which allows the expansion in the number of goods.[17] Ideally the results obtained with the CN classification should be the preferred ones. After all, the aim of the analysis is to compute the welfare gains due to variety growth and the CN system gives the best measure of variety. Unfortunately, once again, the problem is related to the definition of the goods. The first lines of the previous two tables show that even though the data are more disaggregated under the CN system, the number of defined good is lower: 1,079 versus 2,466 obtained with the HS classification. Since many new categories are appearing in the CN classification, I have to define the goods at a more aggregated level to get identification. This causes a large reduction in the number of goods: 1,079 versus 9,644 categories available at the 8-digit level in 2006. And it can well have an impact on the final results because the lambda ratios depend on the estimated elasticities of substitution, which are likely to vary depending on the level of aggregation at which the goods are defined. It is intuitive that goods defined at a more aggregate level will have a lower value of elasticity because consumers are less likely to substitute; this effect can potentially bias the results upwards because less substitutable goods imply higher welfare gains from variety growth. From this point of view, the HS classification is therefore safer because most of the goods are still defined at the most detailed available level.

I implement two strategies to test the impact of the good definition. The first one consists in reducing the number of goods that cannot be estimated by correcting for missing quantities. Many categories have total trade values but they miss the quantity in some years; this implies a loss of data during the estimation because data points for which the unit value cannot be calculated are dropped. Hence it is not possible to estimate many categories at the 6-digit level even though trade values are available for all years. In order to reduce this problem, I fill

[17] BW aggregate all the former Soviet countries in a fictitious country throughout their analysis. They want to rule out any spurious increase in the number of varieties due to the dissolution of the Soviet Union and the creation of other 14 independent nations during 1990 and 1991. This approach is definitely right in the aftermath of the dissolution because there is a mechanical increase in the number of varieties due to new independent countries now exporting the same old Soviet products. But so is it wrong in later years. If, say, Latvia and Kyrgyzstan start exporting a new good in 1998, why should these two goods be considered as the same variety coming from a fictitious post-USSR? Hence there is a clear underestimation of variety growth in later years when this growth was actually higher. In 1993 former Soviet countries exported a total of 7,042 varieties, which become 6,534 if the former Soviet Republics are aggregated in a fictitious post-USSR. Hence the spurious increase in varieties is around 7.7%, which is quite a low number compared to the total 180% growth of Figure 3.1. If the fictitious aggregation is performed throughout all years anyway, the total variety growth from ex-Soviet countries decreases to 147% and the total welfare gains decrease to 1.74% from 2.08%. Hence, even completely shutting down this channel, the picture does not change much. The correct estimate is obviously in the middle, but probably much closer to 2%, given the higher variety growth in later years, which can be safely considered as real. Moreover, given the very conservative approach taken throughout the analysis and the clear underestimation of variety growth due to the HS classification, it is reasonably safe to leave the former Soviet Republics as independent nations.

Table 3.7: The Impact of Good Definition, Harmonized System

Level	Number of goods	Ratio to total trade	Model (1) Sigma	Model (1) Welfare	Model (2) Sigma	Model (2) Welfare	Model (3) Sigma	Model (3) Welfare
Digit 2	96	90.36%	3.42	2.66%	7.18	1.20%	3.38	1.84%
Digit 3	173	90.36%	3.19	2.58%	7.30	1.04%	3.12	1.76%
Digit 4	1099	90.36%	7.12	2.52%	12.79	1.25%	7.44	2.22%
Digit 5	2906	90.36%	8.45	2.39%	17.35*	1.19%	9.16	2.11%
Digit 6	3882	90.36%	9.45	2.49%	17.07*	1.18%	10.58	2.25%
All	2466	100%	8.28	1.85%	16.60*	1.07%	8.78	2.08%

Data: OECD - ITCS - HS 1988. *: one outlier is dropped.

missing quantities with a linear interpolation. The results of this exercise are shown in Table A.1, the total number of goods goes up to 2,935 from 2,466 and, in particular, the number of goods estimated at the 6-digit rises to 2,664 from 2,146. As expected, welfare gains are generally smaller because goods are estimated at a relatively more disaggregated level. This can be rationalized by the higher elasticity estimates: the median elasticity is higher for all model specifications. Nevertheless, interestingly enough, the overall impact is not so big. The welfare gains for the unweighted models are almost exactly the same and the gains for model 1 are even higher, almost equal to the total estimate of model 3. Two opposite effects are likely to be present: higher elasticity estimates on average reduce the welfare gains ('direct effect') but the higher number of goods implies more possible outliers, and, given the non-linearity of the problem, few goods with low elasticity values can drive the total welfare gains upwards ('outlier effect').

A drawback of the previous strategy is that the underlying data are not exactly the same, so the result can be partially driven by that. To overcome this problem and to clearly isolate the effect of the good definition, I propose a second way to solve the identification problem. The idea is quite simple and consists in keeping the categories that can be identified at the 6-digit level only (or 8-digit for the CN); everything that cannot be identified is thrown away and all the goods are therefore defined at the 6-digit level. Then, keeping the underlying data fixed, the level at which the goods are defined is changed, hence isolating the effect of the good definition.

The results of the exercise for the HS are shown in Table 3.7. It easy to note that the total number of goods grows a lot: 3,882 categories can be identified at the 6-digit level (penultimate line in the table). The reason is that in the process of aggregation, many categories that are actually identifiable are joined with those that are not, all together they become varieties of the same (more aggregate) good that is now identifiable. For the data from the HS, the number of categories that is thrown away is not that big, in terms of total trade this is just 9.6%; hence variety growth is estimated using 90.4% of total trade. Ultimately, the aim of the exercise is to quantify the relative change of the welfare gains estimates when the level at which goods are defined is changed, and not the absolute gains. But in order to make the estimates comparable with the results obtained with the total value of trade, the welfare gains are calculated as if the estimated subset of categories represented the total of trade. That is, the log-change ideal weights are recomputed and sum up to one and the ideal import share is not modified, so it

Table 3.8: The Impact of Variety Growth with Long Run Elasticities

Quantity	Statistic	Model (0)	Model (1)	Model (2)	Model (3)
Num. Of Goods		2373	2373	2373	2373
3 digit		67	67	67	67
4 digit		148	148	148	148
5 digit		109	109	109	109
6 digit		2049	2049	2049	2049
Sigma	Mean	14.18*	14.51	14.84*	15.86
	St. err.	0.97	1.25	0.82	1.36
	Percentile 5	1.72	1.41	1.71	1.50
	Median	3.58	3.04	3.64	3.20
	Percentile 95	135.50	135.50	135.50	135.50
Lambda Ratio	Percentile 5	0.681	0.522	0.681	0.584
	Median	0.978	0.973	0.979	0.976
	Percentile 95	1.093	1.116	1.088	1.113
Bias	Estimate	0.997	0.951	0.997	0.961
Welfare Gains	Estimate	0.08%	1.20%	0.07%	0.95%
Soviet Countries	2006	16.92%	15.15%	18.89%	14.43%
	Mean	44.27%	11.54%	50.83%	10.90%

Data: OECD - ITCS - HS 1988.

still represents the total of imports. It is not possible to claim that the results are perfectly comparable because the elimination of some categories can well bias them in either direction, but it makes the comparison more direct and easier.

The results obtained with the preferred specification (model 3) again show that the level at which the goods are defined does not have such a big impact; the welfare gains are quite robust remaining very close to 2%. As one might expect the average value of the estimated sigmas decreases if the goods are defined at a more aggregated level; the average sigma at the 2-digit level is equal to 3.38, down from 10.58 when goods are defined at the 6-digit level. Nevertheless the estimated welfare gains do not grow consequently. The two effects previously described apply here as well. The 'outlier effect' can be seen from the fact that the standard deviation decreases a lot when goods are estimated at a more aggregated level; hence estimates tend to similar values and outliers in either direction are eliminated. Since the entire calculation is highly non-linear, the elimination of few very low values counterbalances the fact that the average sigma is decreasing. The weighting scheme proposed by BW (model 1) instead appears to be more sensitive to the level at which goods are defined; the estimate of welfare gains increases up to 2.7% when the goods are defined at the 2-digit level.[18] This is an extra reason why the modified weighting scheme is preferred to the original one; a stronger robustness to the level at which goods are defined is a desirable property, especially when the goods are defined at different aggregation levels or when the classification used is changed.

[18] A very similar result is found when the data are interpolated (Table 3.B.3). Also in that case the welfare gains for model 3 are quite stable at 2%, while the estimates of model 1 are more sensitive and rise up to 3.1% when the goods are defined at the 3-digit level.

Finally, as it has been shown in the previous section, the estimates of elasticities are sensitive to the time differences used to calculate them. A first attempt to obtain a measure of long run elasticities has been proposed by taking 5-year instead of 1-year time differences, the sensibly higher elasticity estimates have a strong impact on the welfare gains in the case of the HS classification as shown by Table 3.8. In the case of the unweighted models (0 and 2) the welfare gains almost disappear.

3.5.2 Welfare Gains over Time and Robustness to the Classification

In the case of the US, BW find that the welfare gains are bigger over the '70s and '80s compared to the later period. They do not advance precise reasons for this result and just refer to globalization in general. However, the phenomenon of globalization has been constantly present over the past 30 years and, if possible, some of its driving forces have been even stronger in the last decade. Consider the surge of China and, in the case of the US, the trade liberalizations like the North American Free Trade Agreement that came into effect in 1994. Moreover the share of imports in total GDP in the US is bigger in the later period, so one would intuitively expect welfare gains to be bigger in the last decade, given the mechanical effect of the higher weight. It is therefore not entirely clear why variety growth has such a stronger effect over the '70s and '80s compared to the '90s.

A possible concomitant cause has to be searched in the estimation procedure. In fact, BW estimate the first period using a different and more aggregated classification; hence one could expect an upward bias in their results due to the more aggregated level at which goods are defined. As discussed in the previous section this bias is certainly present, but it is probably weaker than expected due to the presence of the two opposite effects: lower sigmas are counterbalanced by a lower standard deviation, and hence less outliers.

A completely different classification brings another complication: variety growth is itself estimated differently due to the restrictions imposed by the classification structure. The previous section has shown the welfare estimates obtained with both the HS and the CN classification. Comparing the two sets of estimates there are two forces going in the same direction: the CN classification has a better measure of variety growth and the final number of estimated goods is lower. The outcome is a significantly higher estimate of the welfare gains in the case of the CN.[19] In order to show a case more similar to BW, I obtain the same estimates for the SITC Rev. 3 classification. In this case total variety growth is likely to be underestimated compared to the HS because the number of available categories is lower, so one would expect lower welfare gains. But, as Table 3.B.2 shows, the final number of estimated goods is also lower and the elasticity estimates are consequently smaller, causing higher welfare gains. Hence the two forces are going in opposite directions and the overall effect cannot be predicted a priori. The results in Table 3.B.2 show that the welfare gains are actually higher for the SITC classification in the

[19] A lower number of goods implies more aggregated goods on average and in principle lower elasticity estimates. However, the elasticity estimates are quite similar in the two cases (lower for the unweighted models but actually higher for the weighted ones, at least for the percentiles displayed), so the higher welfare gains are likely to be driven mainly by the more precise measure of variety growth.

case of the weighted models, in particular for model 1. So it seems that the bias caused by the lower elasticity estimates prevails in this case; and it is particularly strong for model 1, which appears to be more sensitive to that type of bias, as the results in the previous section showed.

The same type of bias can well have played a role in the BW's analysis. In fact the total number of categories available in the first period under the TSUSA classification was 12,347 versus 14,549 available in the later period under the HTS. Variety growth is better measured in the later period but, at the same time, goods and elasticities are estimated at a more aggregated level in the early period. Therefore the two forces go again in opposite directions and it is not possible to establish the direction of the bias a priori. If we are in a case similar to the comparison between the SITC and the HS classifications discussed before, then the welfare gains in the first period are likely to be biased upwards. But it is hard to quantify the magnitude of the bias and it is not possible to clearly establish whether the BW's finding of higher welfare gains in the first period are driven by the classification bias or not.

At this stage, it is interesting to investigate whether the same time path holds for the UK, where the comparison of the welfare estimates over time can be performed more precisely. Some important episodes, like the collapse of the Soviet Union and the implementation of the Single Market Programme, took place in the early '90s, and one would expect to find higher welfare gains during this period compared to the '70s and '80s. To rule out the classification problem, I estimate the welfare gains using the same classification over the whole time frame. I use the SITC Rev. 2 classification, which allows me to go back to the '70s but, being less detailed, it has the likely consequence of underestimating variety growth. Yet, here the aim is to establish the pattern of welfare gains over time so the absolute values of the estimates are of secondary importance. In order to eliminate any other possible disturbance, the elasticities and the number of goods are estimated for the whole time framework and kept constant across the different periods. The total time framework is 1978 to 2006 and is split into two sub-periods, allowing for the break year to vary. The results of this exercise are shown in Table 3.9. Very interestingly, in the case of the UK, the welfare gains from variety growth over the late '70s and '80s are nil, if not negative. Most of the gains accrue in the second period, and the break year that has the biggest effect is 1990, with welfare gains equal to -0.35% between 1978 and 1990 and equal to 2.69% in the later years until 2006. The Berlin wall fell in 1989 and the Soviet Union collapsed in 1991, this is probably more than a coincidence.

3.5.3 The Fall of the Iron Curtain

The previous results demonstrate the presence of welfare gains even in a framework where the change in variety is underestimated due to the restrictions imposed by the HS's structure, which does not allow for the expansion in the number of categories. At this stage, it is interesting to shed some light on the prime causes of these gains, and ultimately to understand where variety growth comes from. In the case of the US, welfare gains are found to be bigger during the '70s and '80s. However it is difficult to understand which are the factors that played the biggest role and BW limit themselves to generically indicate globalization as the main cause. In the case

Table 3.9: The Impact of Variety Growth in the UK, 1978-2006

Quantity	Statistic	Period													
		78-89	89-06	78-90	90-06	78-91	91-06	78-92	92-06	78-93	93-06	78-94	94-06	78-95	95-06
Num. Goods		536		536		536		536		536		536		536	
Sigma	Mean	5.80		5.80		5.80		5.80		5.80		5.80		5.80	
	St. err.	0.70		0.70		0.70		0.70		0.70		0.70		0.70	
	Percentile 5	1.57		1.57		1.57		1.57		1.57		1.57		1.57	
	Median	2.86		2.86		2.86		2.86		2.86		2.86		2.86	
	Percentile 95	10.54		10.54		10.54		10.54		10.54		10.54		10.54	
Lambda Ratio	Percentile 5	0.858	0.710	0.857	0.700	0.775	0.709	0.803	0.715	0.746	0.648	0.783	0.800	0.740	0.811
	Median	0.999	0.986	0.999	0.986	0.999	0.985	0.999	0.988	1.000	0.981	0.994	0.994	0.995	0.995
	Percentile 95	1.131	1.220	1.210	1.199	1.135	1.200	1.175	1.209	1.287	1.148	1.146	1.168	1.262	1.159
Bias	Estimate	1.013	0.908	1.016	0.894	1.018	0.920	1.016	0.939	1.005	0.928	0.995	0.936	0.994	0.940
Welfare	Estimate	-0.30%	2.36%	-0.35%	2.69%	-0.38%	1.88%	-0.32%	1.44%	-0.10%	1.75%	0.11%	1.56%	0.13%	1.52%
	Total	2.05%		2.33%		1.50%		1.11%		1.64%		1.67%		1.64%	
Ideal weight		0.223	0.243	0.218	0.237	0.206	0.225	0.208	0.227	0.214	0.233	0.217	0.236	0.224	0.244

Data: UN Comtrade - SITC Rev. 2. Model 3: err_adj_2, weight_2. Data for the ideal weights from the World Bank World Development Indicators.

of the UK the picture might differ; the increase in trade with China and East Asian countries, which can simplified with the term globalization, have surely a great importance. But this is a phenomenon that slowly increased over the past 30 years while the UK and Europe in general have experienced an important and more sudden event in the early '90s, namely the Fall of the Iron Curtain. This revolutionary event deeply modified the socioeconomic environment in a very short amount amount of time. The previous section already showed that, in the case of the UK, the welfare gains due to variety growth are greater over the '90s, starting from 1990 in particular. Given the coincidence with the timing of the Fall of the Iron Curtain, it is interesting to understand whether the higher welfare gains over the '90s can be actually ascribed to the expansion of trade with the former Soviet countries.

Figure 3.1 already showed quite explicitly the important role played by the former Soviet countries in the growth of varieties imported in the UK. Here I perform a more detailed analysis to precisely estimate the contribution of those countries. A first strategy asks whether the welfare gains from variety growth are bigger for the goods in which ex-Soviet countries have a bigger export share. In order to do so, all the goods are ranked in terms of the ratio between the trade value from those countries and total trade. The goods corresponding to the top percentiles of that ratio are selected and the import bias and the welfare gains are computed. Interestingly enough, the goods for which former Soviet countries have a bigger export share over the period, generate very big welfare gains. The last line of Table 3.10 shows that the top 15% of goods (defined at different disaggregation levels) generate an import bias that is roughly three times bigger than the bias obtained for all of the imported goods. If the total variety growth had matched the growth of the top 308 goods, the welfare gains would have been equal to 6.13% and not just 2.08%. A potential flaw of this analysis is related to the definition of goods. In fact, the aggregation strategy used to overcome the identification problems is affected by the elimination of some of the 6-digit categories; this causes a modification of the definition of the goods that can potentially influence the analysis. In order to rule out all these problems, the goods are once again defined at the 6-digit level only; i.e. all the goods that are not identified at the 6-digit level are eliminated. In this way the comparison with the welfare gains for the total sample is not flawed by the good definition. Results are shown in the first three lines of Table 3.10; even in this case, the goods for which ex-Soviet countries have a bigger export share bring about significantly bigger welfare gains. This result is quite robust to the number of goods included. The picture does not change much if the top 25%, 15% or 5% of goods are considered; in all cases the welfare gains would have been roughly twice as big if the total variety growth had matched the growth of the top goods.

An immediate question arises: which are these goods heavily exported by former Soviet countries that account for so much of the welfare gains? It is hard to show all of them because, even considering the top 168 goods only, the list would be quite long. An alternative way to answer this question is to compare the frequency distribution of the codes of the top 168 goods with the distribution of the total sample of available goods. I implement this simply tabulating the frequency distribution of the first digit of the HS codes, to understand how these goods are distributed across the main categories. The results, reported in Table 3.11, show that the

Table 3.10: Welfare Gains from Soviet Countries, 1988-2006

Level	Number of goods	Top Percentile	Welfare	Welfare Full Sample	Ratio
Digit 6	978	25%	4.00%	2.25%	1.78
Digit 6	553	15%	4.55%	2.25%	2.02
Digit 6	168	5%	4.76%	2.25%	2.12
All	635	25%	3.67%	2.08%	1.76
All	308	15%	6.13%	2.08%	2.94

Data: OECD - ITCS - HS 1988.

Table 3.11: Frequency Distribution of Product Codes

Digit 1	Top 168 Soviet Goods			Total Products		
	Freq.	Percent	Cum.	Freq.	Percent	Cum.
0	7	4.17	4.17	347	6.92	6.92
1	6	3.57	7.74	237	4.72	11.64
2	17	10.12	17.86	762	15.19	26.83
3	5	2.98	20.83	394	7.85	34.68
4	21	12.50	33.33	368	7.34	42.02
5	9	5.36	38.69	501	9.99	52.00
6	36	21.43	60.12	445	8.87	60.87
7	33	19.64	79.76	549	10.94	71.82
8	27	16.07	95.83	1029	20.51	92.33
9	7	4.17	100.00	385	7.67	100.00
Total	168	100		5017	100	

Data: OECD - ITCS - HS 1988.

top 168 products are not concentrated in any particular section of the HS classification; hence the result is not driven by a particular sub-sample of goods. The frequency distribution of the first digit shows that all of the 10 sections are covered and it is more or less in line with the distribution of the first digit of the total available product-categories. A closer look reveals that the top 168 'Soviet products' are slightly less concentrated in the categories starting by 3, which mainly correspond to the Section VI ('Products of the chemical or allied industries') of the HS classification. Instead they display a higher frequency compared to the total sample in the categories starting by 6 and 7; in particular chapters 62 ('Articles of apparel and clothing accessories, not knitted or crocheted'), 64 ('Footwear, gaiters and the like; parts of such articles'), 70 ('Glass and glassware') and 72 ('Iron and steel') of the HS classification.[20]

The previous analysis considers the goods that have a high share of imports from ex-Soviet countries but it cannot rule out the possibility that most of the gains from variety growth might come from other countries that export the same goods. Hence, a better strategy consists in calculating the share of welfare gains that directly comes from those countries. The aim is to calculate the contribution of a country or a set of countries for each good and then sum over all the goods to obtain the total gains from that particular country. In the case of the whole set of

[20] In light of the findings of Goldberg et al. (2010) and Halpern et al. (2011), an interesting avenue for future research is to investigate whether these 'Soviet' goods are more intermediate or final in nature.

ex-Soviet countries, the cost share is calculated as follows:

$$s_{g,sov,t} = \frac{\sum\limits_{k \in I_{sov,t}} p_{gkt} x_{gkt}}{\sum\limits_{c \in I_g} p_{gct} x_{gct}} \tag{3.22}$$

where I_g is the set of common varieties between the starting and the final year of the period, $I_{sov,t}$ is the set of Soviet varieties and t is either 1988 or 2006. In principle one would like to calculate log-change ideal weights for each country and each good but this would require each country to export all of the goods both at the beginning and at the end of the analyzed period. Unfortunately, it is not possible to calculate the contribution of a country if it did not export the good at the beginning of the period because the implied ideal weight would be zero. I therefore use two alternative set of weights. The first set is simply based on the cost shares of each country in 2006; the weight for the ex-Soviet countries for each good is simply their cost share in 2006 divided by the sum of total cost shares (Soviet and non-Soviet). A drawback of these weights is that they are based on the values of the last year only, hence they might overweight countries that have grown faster over the analyzed period. The second set of weights corresponds to the Törnqvist's (1936) weights; that is, in the formula of the log-change ideal weights, an arithmetic mean is substituted for the logarithimic mean as follows:

$$\omega_{g,sov,t} = \frac{\left(s_{g,sov,t} + s_{g,sov,t-1}\right)/2}{\sum\limits_{k \in \{sov, non-sov\}} \left(s_{g,k,t} + s_{g,k,t-1}\right)/2} \tag{3.23}$$

The lambda ratio of each good is now raised to its log-change ideal weight times the Soviet weight just defined. I can therefore calculate the import bias stemming from former Soviet countries only; the total import bias is then simply obtained multiplying the Soviet and non-Soviet import biases: $IB = IB_{sov} \cdot IB_{non-sov}$. The contribution of Soviet countries to total gains is finally calculated as: $SovContr = \ln(IB_{sov})/\ln(IB)$.

This calculation is shown at the bottom of Tables 5, 6 and A.1. Using model 3 and the HS classification the welfare gains from the former Soviet countries account for 10.4% of the total when the 2006 weights are considered, and 7.8% when the mean weights defined in equation (3.23) are used. When the interpolated data are considered (Table 3.B.1) the contribution of ex-Soviet countries becomes 11.6% and 9%, depending on the weights. Finally, looking at the results for the CN classification (Table 3.5) the gains are 13.9% and 11.5% of the total gains. Overall, the contribution of the former Soviet countries accounts for more than 8% of the total welfare gains. The contribution is sizeable considering that the trade value from those countries amounts to 6.4% of total imports in 2006, up from a mere 1.6% in 1988; the average trade share over the 19 years of the analyzed period amounts to 2.9%. The contribution to the welfare gains from variety is therefore roughly three times bigger than the average trade share. As a comparison, the analogous contribution of China accounts for 5% of total gains over the same period. This share is certainly remarkable for a single country; nevertheless, given the exceptional exporting performance that China had over the past years, the results highlight the primary role that also the ex-Soviet bloc has played. Moreover the trade share of China in total imports was 6.8% in 2006 and the average over the period was 3.4%; so, even though the former Soviet countries have overall a lower trade share compared to China, their share in the total gains from variety growth is almost twice as big.

In summary, even from this more detailed analysis, it is clear that the contribution of ex-Soviet countries to the growth in varieties has well exceeded their trade share. The welfare gains from variety growth are cumulative over time so it is very hard to disentangle the single causes that contributed to that. The welfare gains ascribed to the former Soviet countries are obviously affected by many factors that took place over the past twenty years; and isolating the effect of the fall of the Iron Curtain would be very hard, given the profound socioeconomic changes that it caused. A counterfactual world with the Iron Curtain still in place would be even hard to imagine and surely those welfare gains would have not materialized, had the Eastern bloc remained as close as before. The fall of the Iron Curtain was certainly the triggering event for all that followed.

3.6 Conclusions

The present study applies an approach similar to Broda and Weinstein (2006) and Feenstra (1994) to investigate the effects of variety growth in the UK over the period 1978-2006. The methodology is modified by providing a refinement of the measurement error analysis and an improved form of the variance, which is at the basis of the WLS estimator. I perform an extensive robustness analysis to investigate the importance of the various assumptions and to assess the impact on the final estimates. Results show that the elasticity estimates are quite sensitive to the weighting scheme. In particular, the unweighted models provide elasticity values that are significantly higher across the board and, consequently, the estimated welfare gains from variety growth are smaller, almost halved when the HS classification is used. The elasticity estimates are also sensitive to the time differences used to calculate them. I propose a measure of elasticities that are better suited to analyze long run changes by taking 5-year differences instead of 1-year, I get estimates that are sensibly higher and have again a strong impact on the welfare gains. Moreover, the strategy used to define the set of estimated goods has also an impact on the final results; in fact, if goods are defined at a more aggregated level the values of elasticity are smaller and this brings about higher welfare gains. Nonetheless the impact is smaller than one would might expect due to the presence of an opposite effect: at more aggregated levels the standard deviation of the estimates decreases and, given the non-linearity of the problem, a smaller number of goods with very low elasticities reduces the welfare gains. Overall, the model characterized by the proposed weighting scheme has proven to be more robust to the definition of goods and to the classification used.

Moreover, this paper tries to shed light on the sources of variety growth. Conversely to the US, most of the welfare gains from variety growth in the UK accrue over the last two decades and are nil, if not negative, over the late '70s and '80s. In particular the break year that has the largest effect is 1990, one year after the fall of the Berlin Wall. This is not just a coincidence, it turns out that the fall of the Iron Curtain and the expansion of trade with the countries of the former Soviet bloc has a positive effect on the welfare of the UK. The number of varieties coming from those countries increased by almost 180%, while total varieties by a mere 20%. Overall, new varieties account for an increase in welfare equal to 2% of GDP between 1988 and 2006 in the most conservative estimate, which is obtained using the HS classification that is stable over time. When the CN classification is used, the welfare gains are three times bigger. Variety growth coming from former Soviet countries accounts for more than 8% of the total welfare gains. This is a sizeable contribution considering that it is roughly three times bigger than the average trade share of those countries. In comparison the share of China in total gains is 5%, and its average trade share over the period is 3.4%. If the total growth in varieties had matched the growth from ex-Soviet countries, the total welfare gains would have been at least twice as big.

Figures

Figure 3.1: Variety Growth from Ex-Soviet Countries

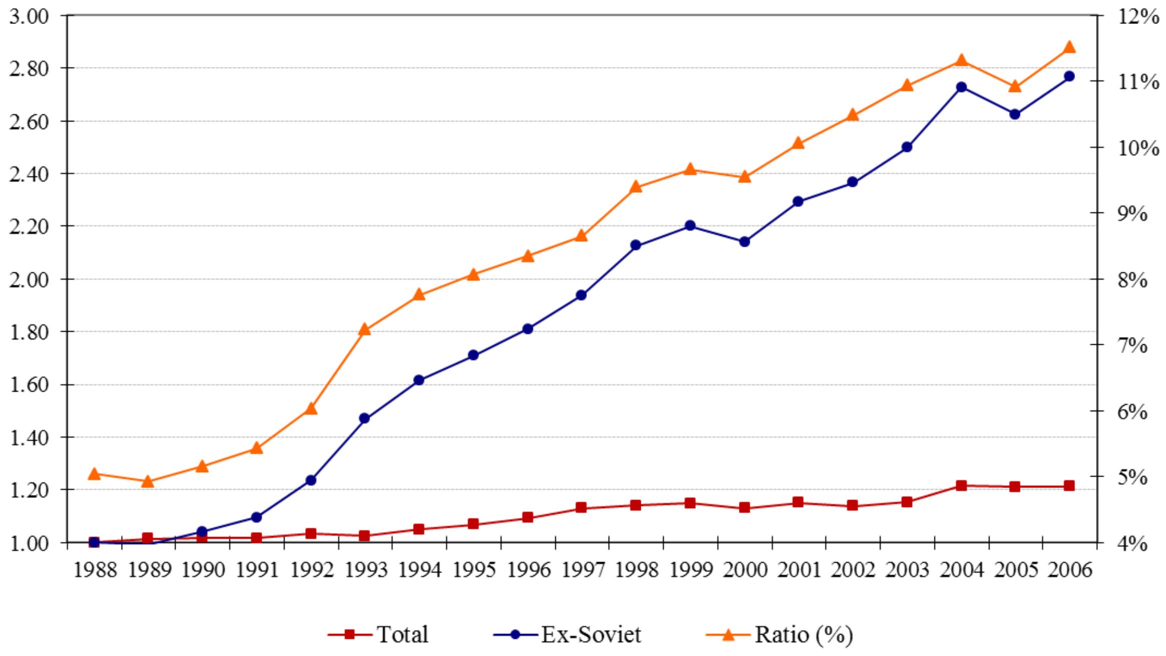

Source: Comtrade and author's calculations.
Note: The left-hand side axis shows the increase in the total number of varieties (Total) and the varieties coming from Ex-Soviet Countries (Ex-Soviet) with their level normalized to 1 in 1988. The right-hand side axis shows the weight of Ex-Soviet varieties in the total (Ratio) over time.

Figure 3.2: Trade Growth from Ex-Soviet Countries

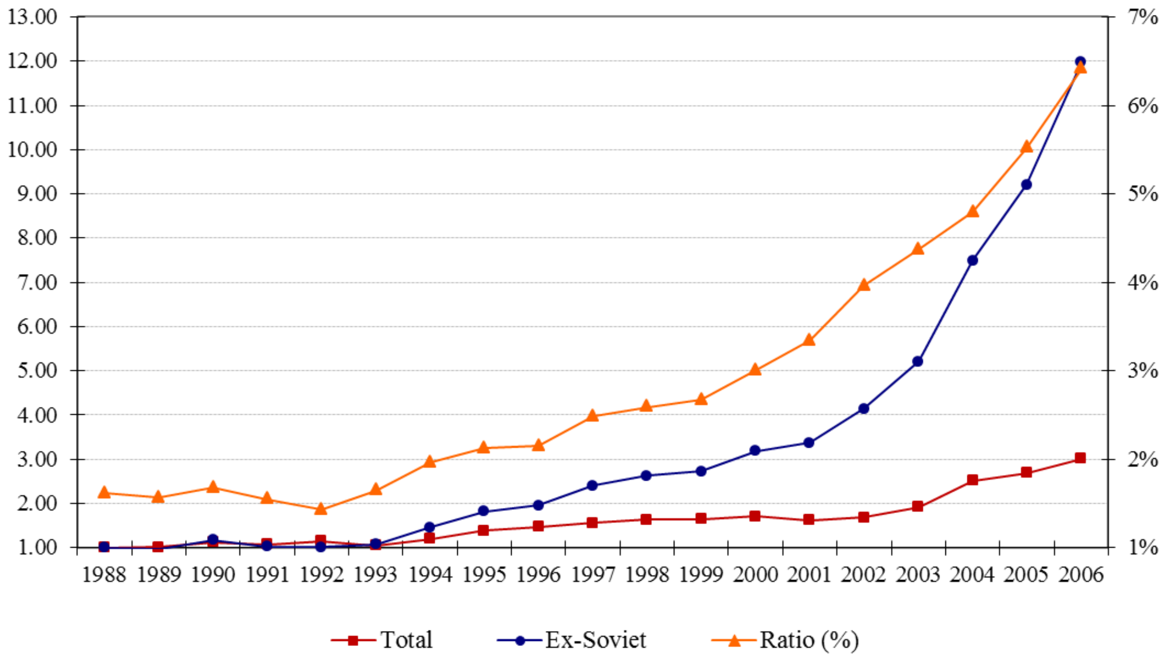

Source: Comtrade and author's calculations.
Note: The figure shows the growth of total imports and imports coming from ex-Soviet countries in the UK. See notes of Figure 3.1.

Figure 3.3: Elasticity Frequency Distributions

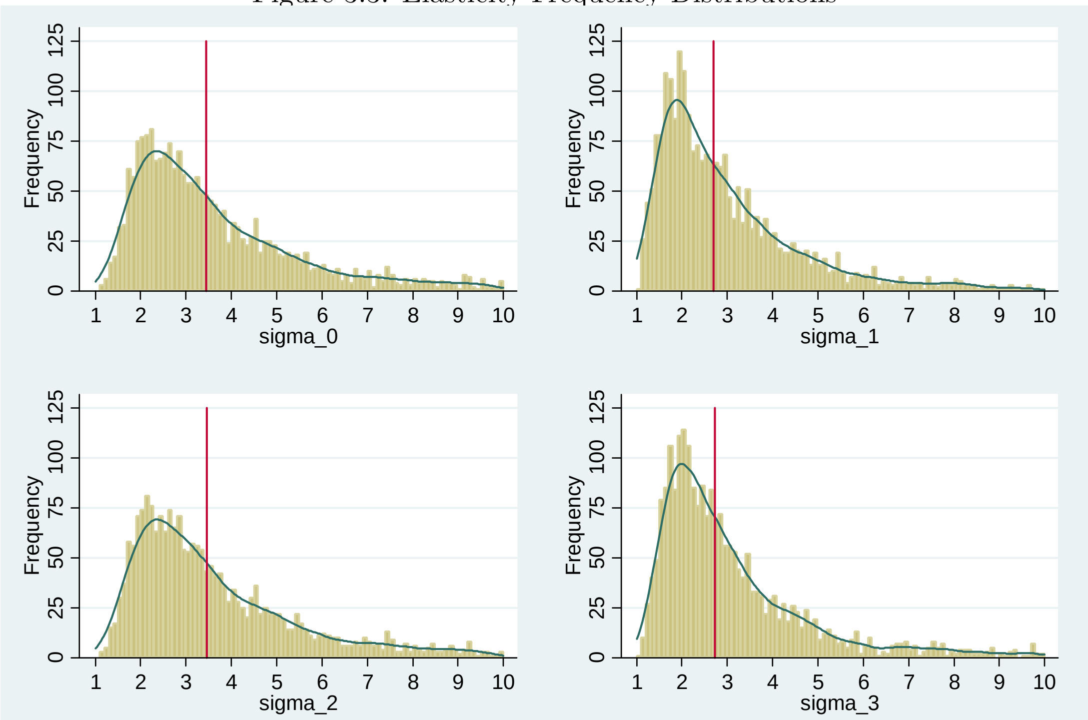

Source: Comtrade and author's calculations.
Note: The figure displays the frequency distribution, the median, and the kernel density of the elasticity estimates according to the four different models. For visual ease, the distributions are truncated at the value of 10.

Figure 3.4: Long Run Elasticity Frequency Distributions

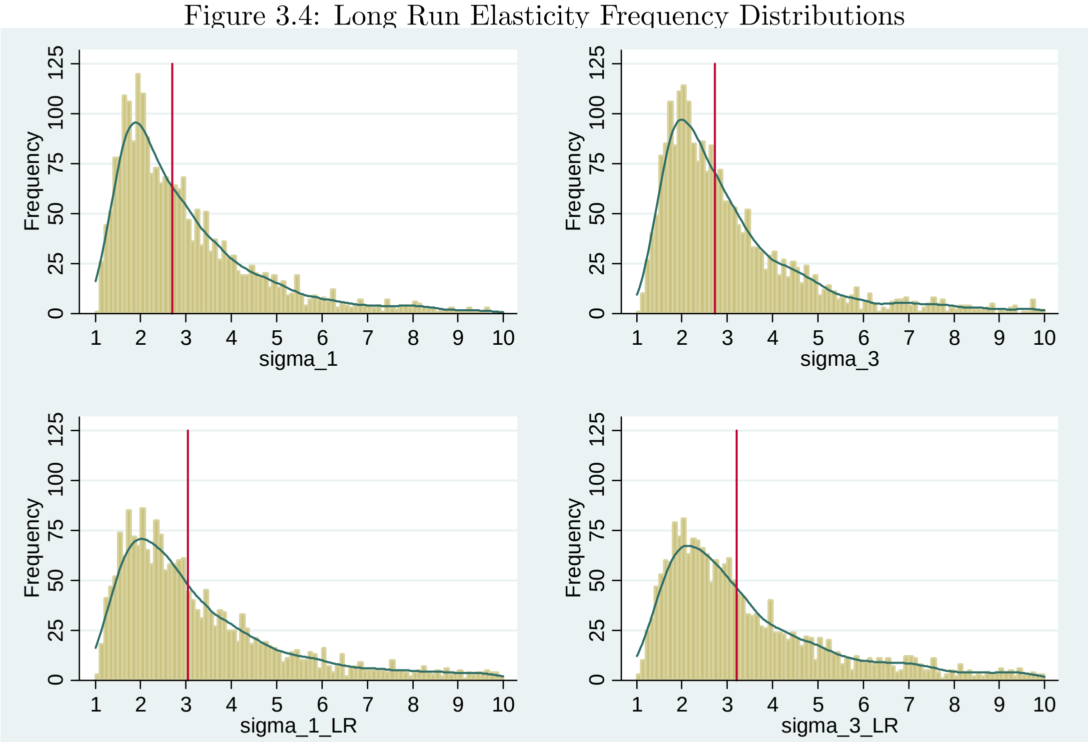

Source: Comtrade and author's calculations.
Note: The top two panels corresponds to the two right hand side panels of Figure 3.3. The bottom two panels display the elasticity estimates obtained with 5-year time differences. See notes in Figure 3.3.

Appendix

3.A Theory

3.A.1 Log-Change Ideal Weights

The log-change ideal weights, ω_{gt}, used to calculate the import bias in Proposition 2 are computed using the good-specific cost shares s_g. The cost share for a particular year is defined as follows:

$$s_{gt} = \frac{\sum\limits_{c \in I_{gt}} p_{gct} x_{gct}}{\sum\limits_{g \in G} \sum\limits_{c \in I_g} p_{gct} x_{gct}} \tag{3.A.1}$$

where G is the set of all goods which remains constant over the whole period, I_g is the set of common varieties between the starting and the final year of the period, $p_{gct} x_{gct}$ is the trade value of a particular variety in year t. Hence the ideal weight is calculated as follows:

$$\omega_{gt} = \frac{(s_{gt} - s_{gt-1}) \, / \, (\ln s_{gt} - \ln s_{gt-1})}{\sum\limits_{g \in G} ((s_{gt} - s_{gt-1}) \, / \, (\ln s_{gt} - \ln s_{gt-1}))} \tag{3.A.2}$$

The log-change ideal weights, ω_t^M, which correspond to the ideal import share used to calculate the welfare gains over the considered periods are defined as follows:

$$\omega_t^M = (s_{Mt} - s_{Mt-1}) \, / \, (\ln s_{Mt} - \ln s_{Mt-1}) \tag{3.A.3}$$

where $s_{Mt} = \frac{\sum_{g \in G} \sum_{c \in I_{gt}} p_{gct} x_{gct}}{GDP_t}$; the numerator of s_{Mt} represents the total goods imports in year t and the denominator is the Gross Domestic Product, both in current US$.

3.A.2 Methodology for the Estimation of the Elasticity of Substitution

The estimation strategy follows Feenstra (1994). By inverting equation (3.2) one can obtain the indirect utility function: $v(p_{gct}, W) = (\Phi_{gt}^M)^{-1} W$, where W is income. Applying Roy's Identity it is possible to work out the import demand equation for each variety of good g, which is defined as follows:

$$x_{gct} = \frac{d_{gct} (p_{gct})^{-\sigma_g} W}{\sum\limits_{c \in I_{gt}} d_{gct} (p_{gct})^{1-\sigma_g}} \tag{3.A.4}$$

From the quantity, the cost share of a particular variety is obtained as follows:

$$s_{gct} = \frac{p_{gct} x_{gct}}{\sum\limits_{c \in I_g} p_{gct} x_{gct}} = \frac{d_{gct} (p_{gct})^{1-\sigma_g}}{\sum\limits_{c \in I_g} d_{gct} (p_{gct})^{1-\sigma_g}} = \frac{\Phi_{gt}^M (I_{gt}, \mathbf{d}_{gt})^{\sigma_g - 1} d_{gct}}{(p_{gct})^{\sigma_g - 1} \lambda_{gt}} \tag{3.A.5}$$

where $\mathbf{d}_{gt}$ is the vector of quality parameters for each country and $I_g = (I_{gt} \cap I_{gt-1})$.

So the import demand equation for each variety of good g can be expressed in terms of shares and changes over time:

$$\Delta \ln s_{gct} = \varphi_{gt} - (\sigma_g - 1)\Delta \ln p_{gct} + \varepsilon_{gct} \tag{3.A.6}$$

where: $\varphi_{gt} = (\sigma_g - 1) \ln P_g^M$ and $\varepsilon_{gct} = \Delta \ln d_{gct}$ is treated as an unobservable random variable, reflecting changes in quality.

Unfortunately it might well be that both $\Delta \ln s_{gct}$ and $\Delta \ln p_{gct}$ are correlated with the error term due to the simultaneous determination of import prices and quantities. So equation (3.A.6) cannot be directly estimates and some assumptions on the supply side of the economy have to be made. Simultaneity bias is corrected by allowing the supply of variety c to vary with the amount of exports, the export supply equation is defined as follows:

$$\Delta \ln p_{gct} = \omega_g \Delta \ln x_{gct} + \Delta \ln v_{gct} \tag{3.A.7}$$

$$x_{gct} = \frac{s_{gct} E_{gt}}{p_{gct}} \tag{3.A.8}$$

$$\Delta \ln p_{gct} = \psi_{gt} + \frac{\omega_g}{1 + \omega_g} \Delta \ln s_{gct} + \delta_{gct} \tag{3.A.9}$$

where: $\psi_{gt} = \omega_g \Delta \ln E_{gt}/(1 + \omega_g)$, E_{gt} is total expenditures on good g, ω_g is the inverse supply elasticity (assumed to be the same across countries) and $\delta_{gct} = \Delta \ln\, v_{gct}/(1 + \omega_g)$ captures any random changes in a technology factor v_{gct}.

The identification strategy relies on the following assumption:

$$E(\varepsilon_{gct}\delta_{gct}) = 0 \tag{3.A.10}$$

This implies that demand and supply equations errors at the variety level are uncorrelated. It is convenient to eliminate φ_{gt} and ψ_{gt} by choosing a reference country k and differencing demand and supply equations, denoted in (3.A.6) and (3.A.9), relative to country k.

$$\Delta^k \ln s_{gct} = -(\sigma_g - 1)\Delta^k \ln p_{gct} + \varepsilon^k_{gct} \tag{3.A.11}$$

$$\Delta^k \ln p_{gct} = \frac{\omega_g}{1 + \omega_g} \Delta^k \ln s_{gct} + \delta^k_{gct} \tag{3.A.12}$$

where $\Delta^k x_{gct} = \Delta x_{gct} - \Delta x_{gkt}$, $\varepsilon^k_{gct} = \varepsilon_{gct} - \varepsilon_{gkt}$ and $\delta^k_{gct} = \delta_{gct} - \delta_{gkt}$. Equation (3.A.12) can be re-written as follows:

$$\Delta^k \ln p_{gct} = \frac{\rho_g}{(\sigma_g - 1)(1 - \rho_g)} \Delta^k \ln s_{gct} + \delta^k_{gct} \tag{3.A.13}$$

where $\rho_g = \omega_g (\sigma_g - 1) / (1 + \omega_g \sigma_g)$ and it satisfies $0 \leq \rho_g \leq (\sigma_g - 1)/\sigma_g < 1$. In order to take advantage of the identification strategy equation (3.A.11) and (3.A.13) are then multiplied together to obtain:

$$(\Delta^k \ln p_{gct})^2 = \theta_1 (\Delta^k \ln s_{gct})^2 + \theta_2 (\Delta^k \ln p_{gct} \Delta^k \ln s_{gct}) + u_{gct} \tag{3.A.14}$$

where: $\theta_1 = \rho_g/(1-\rho_g)(\sigma_g-1)^2$, $\theta_2 = (2\rho_g-1)/(1-\rho_g)(\sigma_g-1)$ and $u_{gct} = \varepsilon^k_{gct}\delta^k_{gct}/(\sigma_g-1)$. The inverse relationship between the thetas, ρ_g and σ_g is given by the following proposition:

Proposition 3. So long as $\theta_1 > 0$, then σ_g and ρ_g are defined as follows:

$$\rho_g = \frac{1}{2} + \left(\frac{1}{4} - \frac{1}{4+(\theta_2^2/\theta_1)}\right)^{1/2} \text{ if } \theta_2 > 0 \tag{3.A.15}$$

$$\rho_g = \frac{1}{2} - \left(\frac{1}{4} - \frac{1}{4+(\theta_2^2/\theta_1)}\right)^{1/2} \text{ if } \theta_2 < 0 \tag{3.A.16}$$

$$\sigma_g = 1 + \left(\frac{2\rho_g - 1}{1-\rho_g}\right)\frac{1}{\theta_2} \text{ in both cases} \tag{3.A.17}$$

If $\theta_1 < 0$, but $\theta_1 > -\theta_2^2/4$, it is still possible to obtain a value for σ_g exceeding unity but $\rho \notin [0,1]$.

3.B Extra Results

Table 3.B.1: The Impact of Variety Growth in the UK, 1988-2006, HS, Interpolated Data

Quantity	Statistic	Model (0)	Model (1)	Model (2)	Model (3)
Num. Of Goods		2935	2935	2935	2935
3 digit		49	49	49	49
4 digit		130	130	130	130
5 digit		92	92	92	92
6 digit		2664	2664	2664	2664
Sigma	Mean	14.18	8.17**	15.27*	9.17
	St. err.	0.99	0.42	0.94	0.49
	Percentile 5	1.71	1.44	1.71	1.53
	Median	3.49	2.81	3.55	2.89
	St. err.	0.05	0.04	0.05	0.03
	Percentile 95	135.48	20.89	135.49	25.18
Lambda Ratio	Percentile 5	0.585	0.467	0.588	0.521
	Median	0.976	0.966	0.977	0.968
	Percentile 95	1.123	1.166	1.119	1.153
Bias	Estimate	0.956	0.923	0.957	0.923
	Conf. Int.	[0.937, 0.968]	[0.854, 0.942]	[0.937, 0.965]	[0.855, 0.949]
Welfare Gains	Estimate	1.08%	1.91%	1.04%	1.93%
	Conf. Int.	[0.78%, 1.55%]	[1.42%, 3.82%]	[0.84%, 1.56%]	[1.24%, 3.80%]
Soviet Countries	2006	11.02%	9.66%	11.13%	11.59%
	Mean	8.60%	7.06%	8.71%	9.04%

Data: OECD - ITCS - HS 1988. *,**: one (77,140) and two (2,385 and 4,246) outliers are excluded, respectively.
See notes in Table 4.

Table 3.B.2: The Impact of Variety Growth in the UK, 1988-2006, SITC

Quantity	Statistic	Model (0)	Model (1)	Model (2)	Model (3)
Num. Of Goods		1921	1921	1921	1921
2 digit		8	8	8	8
3 digit		23	23	23	23
4 digit		335	335	335	335
5 digit		1555	1555	1555	1555
Sigma	Mean	15.00*	8.00	15.21	8.85
	St. err.	0.93	0.77	0.93	0.90
	Percentile 5	1.74	1.38	1.75	1.46
	Median	3.38	2.63	3.40	2.69
	St. err.	0.06	0.04	0.06	0.04
	Percentile 95	135.50	15.75	135.50	17.74
Lambda Ratio	Percentile 5	0.603	0.478	0.603	0.513
	Median	0.979	0.966	0.979	0.967
	Percentile 95	1.156	1.233	1.145	1.209
Bias	Estimate	0.960	0.892	0.961	0.901
	Conf. Int.	[0.918, 0.974]	[0.784, 0.910]	[0.936, 0.972]	[0.861, 0.936]
Welfare Gains	Estimate	0.97%	2.74%	0.96%	2.51%
	Conf. Int.	[0.64%, 2.05%]	[2.23%, 5.95%]	[0.67%, 1.59%]	[1.59%, 3.60%]

Data: UN Comtrade - SITC Rev. 3; *: one outlier (3,637) is excluded. Period: 1988-2006.

Table 3.B.3: The Impact of Good Definition, Harmonized System, Interpolated Data

Level	Number of goods	Ratio to total trade	Model (1) Sigma	Model (1) Welfare	Model (2) Sigma	Model (2) Welfare	Model (3) Sigma	Model (3) Welfare
Digit 2	96	91.47%	3.49	2.66%	6.10	1.33%	3.41	1.70%
Digit 3	173	91.47%	4.36	3.12%	11.89	1.45%	3.38	2.03%
Digit 4	1135	91.47%	7.79	2.99%	11.49	1.47%	8.16	2.05%
Digit 5	3051	91.47%	9.24	2.30%	15.93	1.23%	11.21	2.08%
Digit 6	4136	91.47%	11.40	2.39%	16.38*	1.19%	12.38	2.23%
All	2935	100%	8.17**	1.91%	15.27*	1.04%	9.17	1.93%

Data: OECD - ITCS - HS 1988. *,**: one and two outliers are excluded, respectively.

List of Tables

List of Figures

Bibliography

Abraham, K. G. and Taylor, S. K. (1996). Firms' use of outside contractors: Theory and evidence. *Journal of Labor Economics*, 14(3):394–424.

Abramovsky, L. and Griffith, R. (2006). Outsourcing and offshoring of business services: How important is ICT? *Journal of the European Economic Association*, 4(2-3):594–601.

Abramovsky, L., Griffith, R., and Sako, M. (2004). Offshoring of business services and its impact on the UK economy. Briefing Note 51, Institute for Fiscal Studies.

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

Acemoglu, D. and Guerrieri, V. (2008). Capital deepening and nonbalanced economic growth. *Journal of Political Economy*, 116(3):467–498.

Alonso, R., Dessein, W., and Matouschek, N. (2008). When does coordination require centralization? *American Economic Review*, 98(1):145–79.

Alvarenga, C. and Malmierca, P. (2010). Core competency 2.0: The case for outsourcing supply chain management. Research paper, Accenture.

Amiti, M. and Wei, S.-J. (2009). Service offshoring and productivity: Evidence from the US. *The World Economy*, 32(2):203–220.

Angrist, J. D. and Pischke, J.-S. (2008a). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Angrist, J. D. and Pischke, J.-S. (2008b). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Antràs, P. (2003). Firms, contracts, and trade structure. *The Quarterly Journal of Economics*, 118(4):1375–1418.

Antràs, P. and Helpman, E. (2004). Global sourcing. *Journal of Political Economy*, 112(3):552–580.

Antràs, P. and Helpman, E. (2008). Contractual frictions and global sourcing. In Helpman, E., Marin, D., and Verdier, T., editors, *The Organization of Firms in a Global Economy*, pages 311–339. Harvard University Press, Chicago, IL.

Arkolakis, C. (2010). Market penetration costs and the new consumers margin in international trade. *Journal of Political Economy*, 118(6):1151–1199.

Arkolakis, C. and Muendler, M.-A. (2010). The extensive margin of exporting products: A firm-level analysis. NBER Working Papers 16641, National Bureau of Economic Research, Inc.

Armington, P. S. (1969). A theory of demand for products distinguished by place of production. *International Monetary Fund Staff Papers*, 16(1):159–178.

Atalay, E., Hortaçsu, A., and Syverson, C. (2012). Why do firms own production chains? Mimeo, University of Chicago.

Autor, D. H. and Dorn, D. (2012). The growth of low skill service jobs and the polarization of the U.S. labor market. *The American Economic Review*, Forthcoming.

Bajari, P. and Tadelis, S. (2001). Incentives versus transaction costs: A theory of procurement contracts. *The RAND Journal of Economics*, 32(3):387–407.

Bartel, A. P., Lach, S., and Sicherman, N. (2009). Outsourcing and technological change. Discussion Paper 4678, IZA.

Bartel, A. P., Lach, S., and Sicherman, N. (2012). Technological change and the make-or-buy decision. *Journal of Law, Economics, and Organization.*

Baum, C. F., Schaffer, M. E., and Stillman, S. (2010). Ivreg2: Stata module for extended instrumental variables/2sls and gmm estimation. Statistical Software Components, Boston College Department of Economics.

Baumol, W. J. (1967). Macroeconomics of unbalanced growth: The anatomy of urban crisis. *American Economic Review*, 57(3):415–426.

Becker, G. S. and Murphy, K. M. (1992). The division of labor, coordination costs, and knowledge. *The Quarterly Journal of Economics*, 107(4):1137–1160.

Berlingieri, G. (2013). The rise of business services and aggregate productivity. Mimeo, London School of Economics.

Bernard, A. B., Redding, S., and Schott, P. K. (2011). Multiproduct firms and trade liberalization. *The Quarterly Journal of Economics*, 126(3):1271–1318.

Blonigen, B. A. and Soderbery, A. (2010). Measuring the benefits of foreign product variety with an accurate variety set. *Journal of International Economics*, 82(2):168–180.

Bloom, N., Draca, M., and Van Reenen, J. (2011). Trade induced technical change? the impact of chinese imports on innovation, it and productivity. CEP Discussion Papers dp1000, Centre for Economic Performance, LSE.

Breusch, T. S. and Pagan, A. R. (1979). A simple test for heteroscedasticity and random coefficient variation. *Econometrica*, 47(5):1287–1294.

Broda, C. and Weinstein, D. E. (2006). Globalization and the gains from variety. *The Quarterly Journal of Economics*, 121(2):541–585.

Buera, F. J. and Kaboski, J. P. (2009). Can traditional theories of structural change fit the data? *Journal of the European Economic Association*, 7(2-3):469–477.

Buera, F. J. and Kaboski, J. P. (2012). The rise of the service economy. *American Economic Review*, 102(6):2540–69.

Caliendo, L., Monte, F., and Rossi-Hansberg, E. (2012). The anatomy of french production hierarchies. NBER Working Papers 18259, National Bureau of Economic Research, Inc.

Caliendo, L. and Parro, F. (2012). Estimates of the trade and welfare effects of nafta. NBER Working Papers 18508, National Bureau of Economic Research.

Caliendo, L. and Rossi-Hansberg, E. (2012). The impact of trade on organization and productivity. *The Quarterly Journal of Economics*, 127(3):1393–1467.

Card, D. (2001). Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics*, 19(1):22–64.

Carvalho, V. M. and Gabaix, X. (2013). The great diversification and its undoing. *The American Economic Review*, 103(5):1697–1727.

Caselli, F. and Coleman II, W. J. (2001). The U.S. structural transformation and regional convergence: A reinterpretation. *Journal of Political Economy*, 109(3):584–616.

Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16):386–405.

Corcos, G., Irac, D., Mion, G., and Verdier, T. (2013). The determinants of intrafirm trade: Evidence from french firms. *The Review of Economics and Statistics*, 95(3).

Costinot, A., Oldenski, L., and Rauch, J. (2011). Adaptation and the boundary of multinational firms. *The Review of Economics and Statistics*, 93(1):298–308.

Crémer, J., Garicano, L., and Prat, A. (2007). Language and the theory of the firm. *The Quarterly Journal of Economics*, 122(1):373–407.

Das, S., Roberts, M. J., and Tybout, J. R. (2007). Market entry costs, producer heterogeneity, and export dynamics. *Econometrica*, 75(3):837–873.

De Loecker, J. (2007). Do exports generate higher productivity? evidence from slovenia. *Journal of International Economics*, 73(1):69–98.

Deblaere, J. G. and Osborne, J. D. (2010). Industrialize and innovate. *Outlook*, (1).

Dessein, W. and Santos, T. (2006). Adaptive organizations. *Journal of Political Economy*, 114(5):956–985.

Dewatripont, M. and Tirole, J. (2005). Modes of communication. *Journal of Political Economy*, 113(6):1217–1238.

di Giovanni, J. and Levchenko, A. A. (2010). Putting the parts together: Trade, vertical linkages, and business cycle comovement. *American Economic Journal: Macroeconomics*, 2(3):95–124.

Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *The American Economic Review*, 67(3):297–308.

Duarte, M. and Restuccia, D. (2010). The role of the structural transformation in aggregate productivity. *The Quarterly Journal of Economics*, 125(1):129–173.

Eaton, J., Kortum, S., and Kramarz, F. (2004). Dissecting trade: Firms, industries, and export destinations. *American Economic Review*, 94(2):150–154.

Eaton, J., Kortum, S., and Kramarz, F. (2011). An anatomy of international trade: Evidence from french firms. *Econometrica*, 79(5):1453–1498.

Echevarria, C. (1997). Changes in sectoral composition associated with economic growth. *International Economic Review*, 38(2):431–52.

Fan, J. P. H. and Lang, L. H. P. (2000). The measurement of relatedness: An application to corporate diversification. *Journal of Business*, 73(4):629–660.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 84(1):157–177.

Feenstra, R. C. (2009). Measuring the gains from trade under monopolistic competition. NBER Working Papers 15593, National Bureau of Economic Research, Inc.

Feenstra, R. C. and Weinstein, D. E. (2010). Globalization, markups, and the u.s. price level. NBER Working Papers 15749, National Bureau of Economic Research, Inc.

Fixler, D. J. and Siegel, D. (1999). Outsourcing and productivity growth in services. *Structural Change and Economic Dynamics*, 10(2):177–194.

Foellmi, R. and Zweimüller, J. (2008). Structural change, Engel's consumption cycles and Kaldor's facts of economic growth. *Journal of Monetary Economics*, 55(7):1317–1328.

Foster, L., Haltiwanger, J. C., and Krizan, C. J. (2001). Aggregate productivity growth. lessons from microeconomic evidence. In *New Developments in Productivity Analysis*, NBER Chapters, pages 303–372. National Bureau of Economic Research.

Fuchs, V. R. (1968). *The Service Economy.* NBER Books. National Bureau of Economic Research.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5):874–904.

Garicano, L. and Rossi-Hansberg, E. (2006). Organization and inequality in a knowledge economy. *The Quarterly Journal of Economics*, 121(4):1383–1435.

Garicano, L. and Rossi-Hansberg, E. (2012a). Organizing growth. *Journal of Economic Theory*, 147(2):623–656.

Garicano, L. and Rossi-Hansberg, E. (2012b). Organizing growth. *Journal of Economic Theory*, 147(2):623–656.

Gibbons, R. (2005). Four formal(izable) theories of the firm? *Journal of Economic Behavior & Organization*, 58:200–245.

Goldberg, P. K., Khandelwal, A. K., Pavcnik, N., and Topalova, P. (2010). Imported intermediate inputs and domestic product growth: Evidence from india. *The Quarterly Journal of Economics*, 125(4):1727–1767.

Gollin, D., Parente, S., and Rogerson, R. (2002). The role of agriculture in development. *American Economic Review*, 92(2):160–164.

Goodman, B. and Steadman, R. (2002). Services: Business demand rivals consumer demand in driving job growth. *Monthly Labor Review*, pages 3–16.

Görg, H., Hanley, A., and Strobl, E. (2008). Productivity effects of international outsourcing: evidence from plant-level data. *Canadian Journal of Economics*, 41(2):670–688.

Hall, B. H., Jaffe, A. B., and Trajtenberg, M. (2001). The NBER patent citation data file: Lessons, insights and methodological tools. NBER Working Papers 8498, National Bureau of Economic Research.

Halpern, L., Koren, M., and Szeidl, A. (2011). Imported inputs and productivity. CeFiG Working Papers 8, Center for Firms in the Global Economy.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054.

Hausmann, R. and Hidalgo, C. A. (2011). The network structure of economic output. *Journal of Economic Growth*, 16(4):309–342.

Hayek, F. A. (1945). The use of knowledge in society. *The American Economic Review*, 35(4):519–530.

Helpman, E. (2006). Trade, FDI, and the organization of firms. *Journal of Economic Literature*, 44(3):589–630.

Helpman, E. and Krugman, P. R. (1985). *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy.* MIT Press, Cambridge, MA.

Helpman, E., Melitz, M., and Rubinstein, Y. (2008). Estimating trade flows: Trading partners and trading volumes. *The Quarterly Journal of Economics*, 123(2):441–487.

Herrendorf, B., Herrington, C., and Valentinyi, A. (2013a). Sectoral technology and structural transformation. Mimeo, Arizona State University.

Herrendorf, B., Rogerson, R., and Valentinyi, A. (2013b). Two perspectives on preferences and structural transformation. *The American Economic Review*, Forthcoming.

Holmström, B. R. (1999). The firm as a subeconomy. *Journal of Law, Economics, and Organization*, 15(1):74–102.

Horowitz, K. J. and Planting, M. A. (2006). Concepts and methods of the input-output accounts. Handbook, Bureau of Economic Analysis. Updated in April 2009.

Horvath, M. (1998). Cyclicality and sectoral linkages: Aggregate fluctuations from independent sectoral shocks. *Review of Economic Dynamics*, 1(4):781–808.

Horvath, M. (2000). Sectoral shocks and aggregate fluctuations. *Journal of Monetary Economics*, 45(1):69–106.

Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, 45(3):511–518.

INSEE (2013). DADS: guide méthodologique. Technical report, Direction des Statistiques Démographiques et Sociales, Institut national de la statistique et des études économiques (INSEE).

Jensen, J. B. and Kletzer, L. G. (2010). Measuring tradable services and the task content of offshorable services jobs. In *Labor in the New Economy*, NBER Chapters, pages 309–335. National Bureau of Economic Research.

Johnson, R. C. and Noguera, G. (2012). Accounting for intermediates: Production sharing and trade in value added. *Journal of International Economics*, 86(2):224–236.

Jones, C. I. (2011a). Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics*, 3(2):1–28.

Jones, C. I. (2011b). Misallocation, economic growth, and input-output economics. NBER Working Papers 16742, National Bureau of Economic Research.

Kleibergen, F. and Paap, R. (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1):97–126.

Kongsamut, P., Rebelo, S., and Xie, D. (2001). Beyond balanced growth. *Review of Economic Studies*, 68(4):869–882.

Krugman, P. R. (1979). Increasing returns, monopolistic competition, and international trade. *Journal of International Economics*, 9(4):469–479.

Krugman, P. R. (1980). Scale economies, product differentiation, and the pattern of trade. *American Economic Review*, 70(5):950–959.

Laitner, J. (2000). Structural change and economic growth. *Review of Economic Studies*, 67(3):545–61.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, 70(2):317–341.

Matsuyama, K. (1992). Agricultural productivity, comparative advantage, and economic growth. *Journal of Economic Theory*, 58(2):317–334.

Mayer, T. and Zignago, S. (2011). Notes on cepiis distances measures: The geodist database. Working Papers 2011-25, CEPII research center.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725.

Meyer, P. B. and Osborne, A. M. (2005). Proposed category system for 1960-2000 census occupations. BLS Working Papers 383, U.S. Bureau of Labor Statistics.

Michaels, G. (2007). The division of labour, coordination, and the demand for information processing. CEPR Discussion Papers 6358, Centre for Economic Policy Research.

Milgrom, P. and Roberts, J. (1995). Complementarities and fit strategy, structure, and organizational change in manufacturing. *Journal of Accounting and Economics*, 19(2-3):179–208.

Morales, E., Sheu, G., and Zahler, A. (2011). Gravity and extended gravity: estimating a structural model of export entry. MPRA Paper 30311, University Library of Munich, Germany.

Nelder, J. A. and Mead, R. (1965). A simplex method for function minimization. *The Computer Journal*, 7(4):308–313.

Ngai, L. R. and Pissarides, C. A. (2004). Structural change in a multi-sector model of growth. CEPR Discussion Papers 4763, Centre for Economic Policy Research.

Ngai, L. R. and Pissarides, C. A. (2007). Structural change in a multisector model of growth. *The American Economic Review*, 97(1):429–443.

Ngai, L. R. and Samaniego, R. (2009). Mapping prices into productivity in multisector growth models. *Journal of Economic Growth*, 14(3):183–204.

Nunn, N. (2007). Relationship-specificity, incomplete contracts, and the pattern of trade. *The Quarterly Journal of Economics*, 122(2):569–600.

Office of Management and Budget (1987). *Standard Industrial Classification Manual*. Executive Office of the President of the United States, Government Printing Office, Washington, D.C.

Oulton, N. (2001). Must the growth rate decline? baumol's unbalanced growth revisited. *Oxford Economic Papers*, 53(4):605–27.

Penrose, E. T. (1959). *The Theory of the Growth of the Firm*. John Wiley and Sons, New York.

Prahalad, C. K. and Hamel, G. (1990). The core competence of the corporation. *Harvard Business Review*, 68(3):79–91.

Quinn, J. B. and Hilmer, F. G. (1994). Strategic outsourcing. *Sloan Management Review*, 35(4):43–55.

Rajan, R. G. and Zingales, L. (2001). The firm as a dedicated hierarchy: A theory of the origins and growth of firms. *The Quarterly Journal of Economics*, 116(3):805–851.

Ramondo, N., Rappoport, V., and Ruhl, K. J. (2012). Horizontal vs. vertical FDI: Revisiting evidence from U.S. multinationals. Mimeo, London School of Economics.

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics*, 48(1):7–35.

Restuccia, D., Yang, D. T., and Zhu, X. (2008). Agriculture and aggregate productivity: A quantitative cross-country analysis. *Journal of Monetary Economics*, 55(2):234–250.

Siegel, D. and Griliches, Z. (1992). Purchased services, outsourcing, computers, and productivity in manufacturing. In Griliches, Z., editor, *Output Measurement in the Service Sectors*, pages 429–460. University of Chicago Press, Chicago, IL.

Silverman, B. S. (2002). *Technological Resources and the Logic of Corporate Diversification*. Routledge, London, UK.

Soderbery, A. (2010). Investigating the asymptotic properties of import elasticity estimates. *Economics Letters*, 109(2):57–62.

Spence, M. (1976). Product differentiation and welfare. *The American Economic Review*, 66(2):407–414.

Tadelis, S. and Williamsonn, O. (2012). Transaction cost economics. In Gibbons, R. and Roberts, J., editors, *The Handbook of Organizational Economics*, pages 159–193. Princeton University Press.

ten Raa, T. and Wolff, E. N. (2001). Outsourcing of services and the productivity recovery in U.S. manufacturing in the 1980s and 1990s. *Journal of Productivity Analysis*, 16:149–165.

Törnqvist, L. (1936). The bank of finland consumption price index. *Bank of Finland Monthly Bulletin*, 10:1–8.

Wernerfelt, B. (1984). A resource-based view of the firm. *Strategic Management Journal*, 5:171–180.

Whelan, K. (2002). A guide to U.S. chain aggregated NIPA data. *Review of Income and Wealth*, 48(2):217–233.

White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, 48(4):817–838.

Williamson, O. E. (1967). Hierarchical control and optimum firm size. *Journal of Political Economy*, 75:123–138.

Winter, S. G. (1988). On coase, competence, and the corporation. *Journal of Law, Economics and Organization*, 4(1):163–80.

Yeaple, S. R. (2006). Offshoring, foreign direct investment, and the structure of U.S. trade. *Journal of the European Economic Association*, 4(2-3):602–611.

Young, P. and Triplett, J. E. (1996). The treatment of auxiliary establishments in industry classification systems. International Establishment Frame Conference, Quebec City.

Yuskavage, R. E., Strassner, E. H., and Medeiros, G. W. (2006). Outsourcing and imported services in BEA's industry accounts. Papers, Bureau of Economic Analysis.