

Skill-Biased Structural Change*

Francisco J. Buera

Washington University in St. Louis

Joseph P. Kaboski

University of Notre Dame

Richard Rogerson

Princeton University

Juan I. Vizcaino

University of Nottingham

November 17, 2020

Abstract

Using a broad panel of advanced economies we document that increases in GDP per capita are associated with a systematic shift in the composition of value added to sectors that are intensive in high-skill labor, a process we label as skill-biased structural change. It follows that further development in these economies leads to an increase in the relative demand for skilled labor. We develop a quantitative two-sector model of this process as a laboratory to assess the sources of the rise of the skill premium in the US and a set of ten other advanced economies, over the period 1977 to 2005. For the US, we find that the sector-specific skill neutral component of technical change accounts for 18-24% of the overall increase of the skill premium due to technical change, and that the mechanism through which this component of technical change affects the skill premium is via skill biased structural change.

*We thank four anonymous referees, Dirk Krueger, Daron Acemoglu, David Dorn, Chad Jones, Pete Klenow, Tommaso Porzio as well as seminar and conference participants at the 2015 AEA Meetings, Arizona State University, the Canon Institute for Global Studies, Chicago Fed, Paris School of Economics, Philadelphia Fed, Pittsburgh, Stanford, University of Houston, USC, Wharton, the World Bank, Yonsei University, UCL, the CEPR, and the Institute for New Structural Economics (INSE) for useful comments.

1 Introduction

The substantial increase in the wages of high-skilled workers relative to low-skilled workers is one of the most prominent secular trends in the US and other advanced economies. A large literature that seeks to isolate the underlying driving forces and propagation mechanisms behind this trend has consistently concluded that skill-biased technological change (SBTC) is a quantitatively important driver of this increase.¹ In this paper we argue that a distinct process – which we label skill-biased *structural* change – has also played a quantitatively important role. We use the term skill biased structural change to describe the *systematic* reallocation of sectoral value added shares toward relatively skill intensive industries that accompanies the growth process among advanced economies and that is driven by sector-specific skill-neutral technical change.

The economic intuition behind our finding is simple. If (as we show is indeed the case in the next section) the process of development is systematically associated with a shift in the composition of value added toward sectors that are intensive in high-skill workers, then the relative demand for high-skilled workers will increase, even if development is driven by the skill-neutral component of technical change. This channel is absent in analyses that adopt an aggregate production function, since in that case the skill-neutral component of technical change has no effect on the relative demand for high-skilled workers.

Our argument proceeds in four steps. We first develop a simple general equilibrium model of structural change that incorporates an important role for skill. To best highlight the shift in value added to relatively skill intensive sectors, we study a two-sector model in which the two sectors are distinguished by their intensity of skilled workers in production. We allow for sector-specific technological change, and decompose technological change in each sector into a skill-neutral component and a skill-biased component. The skill-biased component captures technological change that affects relative marginal products holding inputs fixed, while the skill neutral component captures technical change that affects the amount of output holding inputs

¹Important early contributions to the literature on the skill premium that stress skill-biased technical change include Katz and Murphy (1992), Bound and Johnson (1992), Murphy and Welch (1992), Berman et al. (1994) and Berman et al. (1998). This is not to say that SBTC is the *only* factor at work, as the literature has also highlighted the effect of other factors on overall wage inequality. For example, DiNardo et al. (1996) argue that labor market institutions such as minimum wages and unionization have played an important role in shaping wage inequality overall, Feenstra and Hanson (1999) emphasize the role of offshoring, and Autor et al. (2013) emphasize the role of trade.

fixed. This decomposition is of interest precisely because of the fact that in one-sector models the skill premium depends only on the skill-biased component and is independent of whether total output is affected. We derive a log linear approximation for a special case of our model and show that changes in the sector-specific skill-neutral component of technical change can impact the skill premium, and that the lone mechanism through which this happens is by reallocating activity across sectors that differ in their skill intensity.

In the second step, we use the model to study the evolution of the skill premium in the US economy between 1977 and 2005. We assume that the only exogenous driving forces are technical change and changes in the relative supply of skilled workers. We measure the change in the relative supply of skill directly from the data. We show how the model can be used to infer preference parameters and the components of technical change using data on the growth in aggregate output, relative sectoral prices and the distribution of sectoral value added, the skill premium and changes in sectoral factor shares. Importantly, our calibrated model perfectly matches the observed increase in the skill premium, which rises from 1.33 to 1.88 between 1977 and 2005.

In the third step we use our model to decompose the overall increase of the skill premium into four components: one due to the change in the relative supply of high-skill workers, a second due to skill-biased technical change, a third due to sector-specific skill-neutral change, and a fourth term that represents the interaction between the two types of technical change. If there had been no technical change, our model predicts that the increase in the relative supply of high-skill workers would have lowered the skill premium to 0.87, a drop of 46 percentage points. It follows that technical change created an increase in the skill premium of 101 percentage points. In our benchmark specification, between 18 and 24 percent of this increase comes from changes in the sector-specific skill-neutral component of technological change.²

The fourth and final step quantifies the mechanism through which the sector-specific skill-neutral component of technical change affects the skill premium. We show that this component drives the rise in the size of the high skill-intensive sector; in fact, in its absence, the value added share of the skill-intensive sector would have decreased modestly. The sector-specific skill-neutral component of technical change is also the dominant source of increases in output in our model. We conclude that systematic changes in the composition of value added associated with the process

²The range of estimates reflects the effect of varying the allocation of the interaction term.

of development are an important mechanism in accounting for the rise in the skill premium.

To assess the importance of skill-biased structural change more broadly we repeat the analysis for a set of ten other OECD countries. While the contribution of the sector-specific skill-neutral component of technical change varies across countries, ranging from around 15 percent to almost 50 percent, the median for this sample is 23 percent, very much in line with our estimates for the US.

Our paper is related to many others in two large and distinct literatures, one on SBTC and the skill premium and the other on structural transformation. Important early contributions to the literature on the skill premium include Katz and Murphy (1992), Bound and Johnson (1992), Murphy and Welch (1992), Berman et al. (1994) and Berman et al. (1998). These papers emphasized the role of SBTC because the increase in the skill premium occurred despite a large increase in the relative supply of high-skill workers, and they could not identify other factors that would lead to a large increase in the relative demand for high-skill workers. In particular, while each of them noted compositional changes in demand as a potentially important factor, none of them found this channel to be of first order importance.

Relative to this literature, our contribution is fourfold. First, we analyze the evolution of the skill premium in general equilibrium in a multi-sector economy. Second, we decompose technological change into sector-specific skill-biased and skill-neutral components and assess the contribution of each component to the evolution of the skill premium. Third, we find a large role for the skill-neutral component of technical change and show that the key mechanism through which it affects the skill premium is via structural change. Fourth, we link the driving forces of this structural change to the process of development. In Section 6, we detail the key reasons that our model based approach leads to a different conclusion than the shift-share approach followed by Katz and Murphy (1992).

An early contribution in the second literature is Baumol (1967), with more recent contributions by Kongsamut et al. (2001) and Ngai and Pissarides (2007). (See Herrendorf et al. (2014) for a recent overview.) Relative to this literature our main contribution is to introduce heterogeneity in worker skill levels into the analysis and to organize industries by skill intensity rather than broad sectors. Caselli and Coleman (2001) is an early paper linking structural transformation and human capital. Differently than us, they focus on the movement of resources out of agriculture and

into non-agriculture, and assume that the non-agricultural sector uses only skilled labor.

Two closely related papers to ours are Buera and Kaboski (2012) and Leonardi (2015). Buera and Kaboski (2012) also study the interaction between development and the demand for skill both empirically and theoretically, but their primary theoretical contribution is conceptual, building a somewhat abstract model to illustrate the mechanism. Relative to them we make three contributions. First, we document the empirical patterns for a larger set of countries and along additional dimensions. Second, we develop a model that nests the benchmark quantitative models used to study the skill premium and structural change in isolation. Their model did not include sector-specific technical progress, the driving force behind structural change in our model. Third, we use the model to quantitatively assess the contribution of the sector-specific, skill-neutral component of technical change.

Leonardi (2015) also considers how changes in primitives are propagated via structural change to generate changes in the skill premium. But whereas we focus on the effect of the skill-neutral component of technical change and find significant effects, he focuses on changes in educational attainment and finds relatively small effects. In Section 6 we show that our model predicts similarly small effects for the change that he focuses on.

Our model structure is broadly similar to that of Acemoglu and Guerrieri (2008). Like us, they study the relationship between development and structural change in a model that features heterogeneity in factor intensities across sectors. But whereas we focus on differential intensities of human capital, they focus on differential intensities of physical capital. Cravino and Sotelo (2019) use a framework similar to ours to show how reductions in trade costs affect demand composition and the demand for skill. Cerina et al. (2018) incorporates skill-biased technical change into a model of structural change. Their focus is on labor market polarization and the role of gender differences.³

An outline of the paper follows. Section 2 documents the prevalence of skill biased structural change in a panel of advanced economies, and evidence suggestive of the mechanisms that drive it. Section 3 presents our general equilibrium model

³Ngai and Petrongolo (2014) use a similar framework to show that compositional changes in value added associated with development can explain part of the decrease in the gender wage gap that has occurred in the US over time.

and characterizes the equilibrium. Section 4 shows how the model can be used to account for the evolution of the US economy over the period 1977 to 2005, and in particular how the data can be used to infer preference parameters and the process of technical change. Section 5 presents our main results about the role of different driving forces on the evolution of the skill premium, and the relation between the skill-neutral component of technical change and structural change. Section 6 discusses the relationship of our results to earlier results in the literature. Section 7 reports results for several sensitivity exercises. Section 8 extends the analysis to a set of nine other countries and Section 9 concludes.

2 Skill-Biased Structural Change: Empirics

This section establishes skill-biased structural change as a robust feature of the data for advanced economies. Using data for a broad panel of advanced economies, we document a strong positive correlation in the time series between the level of development in an economy, as measured by GDP per capita, and the share of economic activity accounted for by the relatively skill-intensive sector. The traditional structural change literature documents patterns both in terms of output value added shares as well as employment shares. Similarly, we find that this pattern holds whether we measure the size of the skill-intensive sector in terms of its output value-added share or its share of overall labor compensation. This relationship is robust across countries, and in particular, the experience of the US is very similar to the average pattern found in the data.

The structural change literature emphasizes two key mechanisms: relative price effects and income effects. We document that both of these mechanisms seem relevant for understanding the increasing size of the skill-intensive sector. First, using cross-country data we document a strong positive correlation in the time series between the level of development and the price of the skill-intensive sector relative to other goods and services. Once again, this relationship is robust across countries and the experience of the US is similar to the average pattern.

Second, to document the potential importance of income effects we supplement the aggregate time series panel analysis with evidence on cross-sectional expenditure shares in the US economy. We show that the expenditure of higher income households contains a higher share of value added from the skill-intensive sector.

2.1 Data Sources

The facts reported in this section are based on several datasets. Sectoral value-added shares, sectoral compensation shares and relative sectoral prices come from the EUKLEMS Database (“Basic Table”).⁴ These data exist in comparable form for a panel of countries over the years 1970-2005. The sectoral data are available at roughly the 1 to 2-digit industry level. Our focus is on advanced economies’ growth experience, so following Buera and Kaboski (2012), we focus on the 15 countries with income per capita of at least 9,200 Gheary-Khamis 1990 international dollars at the beginning of the panel in 1970.⁵ Cross-country data for real (chain-weighted) GDP per capita data is from the Penn World Tables 9.0.

Labor compensation data come from the EUKLEMS Labour Input Data. This dataset reports the share of sectoral labor compensation that goes to different skill groups. We define the high-skill group to be those with a college degree or more and define the low-skill group to be all other workers.

Our analysis of US micro data is based on the Consumer Expenditure Survey (CEX), a cross-section data set on household expenditure. This dataset reports household expenditure on final expenditure categories and not value-added categories. To create consistent measures we map household expenditure data through the input-output system to determine the consumption shares of value added. We briefly sketch the steps of this procedure here, and provide more details in the Online Appendix.

We start with the household level CEX data for the United States from 2012 and clean the data as in Aguiar and Bils (2015). We adapt a Bureau of Labor Statistics mapping from disaggregated CEX categories to 76 NIPA Personal Consumption Expenditure (PCE) categories and then utilize a Bureau of Economic Analysis (BEA) mapping of these 76 PCE categories to 69 input-output industries that properly attributes the components going to distribution margins (disaggregated transportation, retail, and wholesale categories). Using the 2012 BEA input-output matrices, we can then infer the quantity of value added of each industry embodied in the CEX expenditures. In our empirical work we restrict ourselves to the primary interview sample, respondents age 24-65 with complete income records, and each observation

⁴See O’Mahony and Timmer (2009).

⁵These countries are Australia, Austria, Belgium, Denmark, France, Germany, Greece, Ireland, Italy, Japan, the Netherlands, Spain, Sweden, the United Kingdom, and the United States. We exclude Luxembourg given its small size. The U.S. data for value added go back to only 1977, while the Japan data go back to only 1973.

is a household-quarter observation.

2.2 Defining the Skill-Intensive Sector

Both our empirical and theoretical analysis will focus on an aggregation of sectors into two bins based on skill intensity: a high-skill intensive sector and a low-skill intensive sector. An important first step is to assign individual sectors to these two broad categories. Our primary metric for assessing skill intensity is the share of total sectoral labor compensation that goes to high-skill workers.⁶ Importantly, the ranking of sectors via this criterion is quite stable over time, so that relative skill intensity can safely be viewed as fixed characteristic of a sector over the period of our analysis.⁷

Creating a binary characterization requires that one adopts a boundary between the high- and low-skill intensive sectors. The benchmark results that we report in this paper are based on defining the skill-intensive sector as consisting of Education, Renting of Machinery and Equipment and Other Business Activities, Financial Intermediation, and Health and Social Work. These four sectors have both the highest average value for their high-skill compensation shares as well as the highest average ranking in the distribution of high-skill compensation shares. In terms of average ranks over the 1977-2005 period, the above four sectors have values of 1.6, 2.4, 3.3 and 4.5 respectively. The next two highest are Chemicals and Chemical Products (5.1) and Real Estate (5.8). In terms of average high-skill compensation shares our four sectors have values of 0.753, 0.528, 0.506, and 0.466. The next two highest are Chemicals and Chemical Products (0.446), and Real Estate (0.434). No other sector has an average above 0.40.

The Real Estate sector merits some discussion. Although this sector ranks quite highly, we have chosen to exclude it from our benchmark definition of the skill-intensive sector. Because this sector has very little employment, its assignment is effectively inconsequential from the perspective of labor variables. But this sector has increased in terms of its value added share since 1977, so including it in the skill-intensive sector would raise the measured increase in the value-added share of the skill-intensive sector. This would serve to increase the size of the effects that

⁶We have also considered the share of total hours or employment accounted for high-skill workers. These alternatives are all highly correlated with our baseline metric based on compensation shares and so do not suggest an alternative ranking. See the Online Appendix for more details.

⁷See the Online Appendix for more details on this.

we estimate, but we feel that this effect is somewhat misleading. For this reason we have decided to err on the side of being conservative and not include Real Estate in our benchmark specification. Because Real Estate and Chemicals and Chemical Products are similarly ranked, we have also chosen to exclude both in our benchmark specification.

However, to assess the possibility that our results are influenced by where we draw the boundary between the two sectors, we have also done our analysis using two more expansive definitions of the skill-intensive sector, one that includes Real Estate and Chemical and Chemical Products, and another that further includes the next two highest ranking sectors, Electrical and Optical Equipment (average rank 7.3) and Public Administration and Defense (average rank 8.1). Using either of these more expansive definition does not affect our main message. Results are included in the Online Appendix.


While our empirical analysis of aggregate time series data is closely related to that in Buera and Kaboski (2012) there is a key difference. They divided industries within the service sector into two mutually exclusive groups: a high-skill intensive group and a low-skill intensive group, and show that whereas the value added share of the high-skill intensive group rose substantially between 1950 and 2000, the value-added share of the low-skill intensive group actually fell over the same time period. In contrast to them, we split the entire economy into a high-skill intensive group and a low-skill intensive group, and not just the service sector. While in our benchmark specification the skill-intensive sector consists exclusively of service sectors, our more expansive definitions also include some goods producing sectors. But importantly, to assess how structural change affects the aggregate demand for skill, one must include the contribution of all sectors and not just those within services. Another difference from Buera and Kaboski (2012) is that we also report cross-sectional micro evidence on income effects.

2.3 Skill-Biased Structural Change

In this subsection we document the phenomenon of skill-biased structural change, i.e., the systematic increase in the relative size of the skill-intensive sector that accompanies development.

Figure 1 shows the relationship between development, as proxied by real GDP per

Figure 1: Structural Change by Skill Intensity and Economic Development.



capita and the rise of the skill-intensive sector.⁸ We use two different measures for the size of the skill-intensive sector: its share of total labor compensation, and its share of total value added. Labor compensation is more relevant from the perspective of labor demand, but value added is the more typical metric for theories of structural change. We include country-level fixed effects with the US being the excluded country. The left panel of Figure 1 shows the relationship using labor compensation, while the right panel shows the relationship using value added. The small squares show the relationship for countries other than the US, and the larger circles represent data for the US.

Both panels lead to the same conclusion: the relative size of the skill-intensive sector increases with log GDP per capita, with highly significant (at 0.1 percent levels) semi-elasticities of 0.21 and 0.15 respectively.⁹ The regression line implies an increase of roughly 30 percentage points of labor compensation and 20 percentage

⁸The analogous figures for the more expansive definition of the skill-intensive sector are contained in the Online Appendix.

⁹The R^2 values for these regressions are 0.93 and 0.92. If we exclude log GDP per capita and only have fixed effects the R^2 values are 0.48 and 0.49, indicating the time series variation in GDP per capita accounts for a large part of the time series variation in the size of the high-skill sector.

points of value added, as we move from a GDP per capita of 10,000 to 40,000 (in 2005 PPP terms).¹⁰ Moreover, we see that the relationship found in the US data is quite similar to the overall relationship. Indeed, the tight relationship suggests that from the perspective of time series changes, cross-country differences in the details for funding of education or health, for example, are second order relative to the income per capita relationship in terms of their effects. (Recall that we have removed country fixed effects in Figure 1.) In sum, the tendency for economic activity to move toward skill-intensive industries as an economy develops is a robust pattern in the cross-country data.

2.4 Structural Change Mechanisms

One common explanation for structural change is changes in relative prices (see, for example, Baumol, 1967; Ngai and Pissarides, 2007). Using value added price indices from the EUKLEMS Database, we examine the correlation between changes in the relative price of the skill-intensive sector and the changes in its value added share that accompanies the process of development.¹¹ Figure 2 is analogous to Figure 1, but plots the value added price index of the high skill-intensive sector relative to the low skill-intensive sector.¹² We have again taken out country fixed effects, and have normalized the relative price indices to 100 in 1995. As before, the larger circles represent the U.S. data.

Figure 2 reveals a strong positive relationship between the relative price of the skill-intensive sector and development.¹³ In this case the relationship in the US data is a bit steeper than in the overall data set, but the strong relationship exists even abstracting from the US. We conclude that changes in relative prices are another robust feature of the structural transformation process involving the movement of activity toward the skill-intensive sector.

A second common explanation for structural change is income effects associated with non-homothetic preferences (see, for example, Kongsamut et al., 2001). With


¹⁰The fact that the semi-elasticity for compensation is significantly higher than for value added will be relevant when we compare our findings with those of Katz and Murphy (1992) later in the paper.

¹¹We construct sector-level aggregate indices as chain-weighted Fisher price indices of the price indices for individual industries. Calculation details are available in the Online Appendix.

¹²The analogous figure for the more expansive definition of the high-skill sector is contained in the Online Appendix.

¹³The R^2 for this regression is 0.76. If we exclude log GDP per capita the R^2 is only 0.10.

Figure 2: Relative Price of Skill-intensive Sector and Economic Development.



this in mind it is of interest to ask whether the output of the skill-intensive sector is a luxury good, i.e., has an income elasticity that exceeds one. To pursue this we examine the relationship between the skill intensity of value-added consumption and income in the Consumer Expenditure Survey (CEX).¹⁴ To the extent that all households face the same prices at a given point in time and have common preferences (or at least preferences that are not directly correlated with income), the cross-sectional expenditure patterns within a country abstract from the relative price relationship in Figure 2 and allow us to focus on the effect of income holding prices constant.

Having constructed household level value added consumption expenditure shares as noted earlier, we regress this share on household observables, most importantly income or education, and potentially a host of other household level controls. Our analysis is similar in spirit to that in Aguiar and Bills (2015) with two exceptions.¹⁵ First, they do not consider our two-sector skill-intensity aggregation, and second,

¹⁴Leonardi (2015) carries out a closely related exercise and concludes that higher income and more educated individuals have higher expenditure shares on final expenditure categories that rely more on high-skill workers, even when taking intermediate input use into account.

¹⁵This exercise is also related to the analysis in Leonardi (2015).

they study final expenditure elasticities whereas we consider value added expenditure elasticities.

Table 1 presents results for regressions of the total share of expenditures that is derived from the skill-intensive sector. The first column presents results from an OLS regression on log after tax income and a set of demographic controls, including age, age squared, dummies for sex, race, state, urban, and month, and values capturing household composition (number of boys aged 2-16, number of girls aged 2-16, number of men over 16, number of women over 16 years, and number of children less than 2 years). The coefficient on log income in the first column indicates that the semi-elasticity of the skill-intensive share of value added embodied in expenditures is 0.030. The second column replaces log income with the log of total expenditures, and finds a larger semi-elasticity of 0.050.¹⁶

Table 1

Household Skill-Intensive Expenditure Share vs. Income or Total Expenditures

	OLS	OLS	IV	IV	OLS
Ln Income	0.030*** (0.001)	-	0.054*** (0.002)	-	-
Ln Expenditures	-	0.050*** (0.002)	-	0.081*** (0.002)	-
High-skill Head	-	-	-	-	0.047*** (0.002)
R^2	0.19	0.22	0.12	0.16	0.18
Observations	13,144	13,210	13,144	13,210	4,056

¹ *** indicate significance at the 1 percent level.

² Standard errors are in parenthesis. Controls include: age; age squared; dummies for sex, race, state, urban, and month; number of boys (2-16 year); number of girls (2-16 years); number of men (over 16 years); number of women (over 16 years); and number of infants (less than 2 years). High-skill is defined as 16 years of schooling attained, while low-skill is defined as 12 years attained. Sample includes households with heads aged 25-64 and complete income data.

Both income and expenditure are certainly subject to measurement error, and even if properly measured, income would only proxy for permanent income, leading to a likely attenuation bias. The third and fourth columns attempt to alleviate this measurement error by instrumenting for log income or log expenditures, respectively,

¹⁶The larger coefficient for expenditures may be driven by certain lumpy expenditures like higher educational expenses and car purchases driving both up in particular months. We nonetheless report these coefficients for the sake of completeness.

using the years of schooling attained by the head of household. Instrumenting for income in this fashion increases the coefficient almost two-fold to 0.054. Likewise, instrumenting for log total expenditures increases the coefficient by about 60 percent to 0.081.

The last column uses education as a direct regressor, replacing log income or log expenditures with a dummy for whether the head of household is high skill or not. Here high-skill is defined as having exactly 16 years of education, while low-skill is defined as having exactly 12 years. (The rest of the households are dropped, leading to the smaller sample size.) The coefficient indicates that the skill-intensive share of value added embodied in expenditures is 4.7 percentage points higher in households with a high-skill head.

We have examined the robustness of the results in Table 1 along various dimensions. Table 1 uses “quarterly” expenditures of the household across the three months they are surveyed, but if we use the monthly data directly, we find nearly identical results. Dropping demographic controls increases the sample by about twenty percent, but again the coefficients are essentially unchanged and highly significant. By defining high-skill as those with at least 16 years of education, and low skill as those with less than 16 years of education, we expand the sample somewhat; the raw coefficient is slightly smaller but not dramatically so (0.029 rather than 0.047). The coefficient remains highly significant. We also examined the diary sample, a smaller sample with a survey that focuses on higher frequency expenditures. In the diary data, we estimate the same coefficient for expenditures but the coefficients on income (0.023 vs. 0.030) and high-skill head (0.022 vs. 0.045) are slightly smaller.¹⁷

Recalling that the aggregate time series data in Figure 1 implied a coefficient of 0.17 on log GDP per capita without controlling for changes in relative prices, the instrumented expenditure coefficient of 0.081 suggests that a significant part of the aggregate time series effect may be driven by income effects. We therefore take this as evidence that, in addition to relative prices, non-homotheticities may also play a role in accounting for the observed pattern of skill-biased structural change

Lastly, we note an important limitation in directly applying the micro elasticity as an income effect. Because the CEX captures only out-of-pocket expenditures, it underestimates the true consumption of certain goods like health care (a substantial

¹⁷Average monthly expenditures in the diary survey are less than ten percent of the average monthly expenditures in the interview survey.

share of which is paid by employers for working individuals and by the government for those on Medicare), and education (a substantial share of which is paid by government).¹⁸ This caution notwithstanding, we will use our estimated elasticity of 0.81 in our calibration exercise as a way to discipline the relative importance of income and relative price effects.¹⁹

2.5 Summary

We have documented a robust relationship in the time series data for advanced economies regarding the systematic movement of activity into the skill-intensive sector associated with the process of development. We refer to this process as skill-biased structural change, so as to emphasize both its connection to the traditional characterization of structural change and the special role of skill intensity. This relationship is remarkably stable across advanced economies, thus suggesting that it is explained by some economic forces that are robustly associated with development, with country specific tax and financing systems not playing a central role in explaining the time series changes.

The traditional structural change literature emphasizes the role of both income and relative price changes as drivers of structural change. We have presented evidence that both of these effects seem relevant in the context of skill-biased structural change as well.

3 Theoretical Framework

Our analysis emphasizes how intratemporal equilibrium allocations are affected by changes in the economic environment that operate through changes in income and relative prices. To capture these interactions in the simplest possible setting, we

¹⁸The estimated income semi-elasticity of the share of out-of-pocket insurance is actually significantly negative in the CEX data although overall insurance consumption is certainly positive. Similarly, although the expenditure share income semi-elasticity of higher education is positive, it is likely understated. Finally, the lack of primary and tertiary expenditures may actually be overstated in the CEX data because it neglects public expenditures, but we conjecture that this relationship is small relative to the higher education relationship.

¹⁹Boppart (2014) also used micro data to discipline these effects, though he used a different two-sector aggregation and also studied final expenditure shares rather than value added shares. Differently from him we will simply use a reduced form elasticity to calibrate our model whereas he used micro data to estimate structural preference parameters.

adopt a static, closed economy model with labor as the only factor of production.²⁰ Our model is essentially a two-sector version of a standard structural transformation model extended to allow for two labor inputs that are distinguished by skill. In this section, we describe the economy and its equilibrium at a point in time; and derive analytic expressions that capture the key economic mechanism at work in our model that connects technical change, structural change, and the skill premium.

3.1 Model

There is a unit measure of households. A fraction f are high-skill, and the remaining fraction $1 - f$ are low-skill. All households have identical preferences defined over two commodities. In our quantitative analysis these two commodities will be connected to the low- and high-skill intensive sectors defined in the previous section. In our benchmark specification all of the high-skill intensive sectors are services and all goods sectors are in the low-skill intensive sector. It is notationally convenient to label the two commodities as goods and services even though what we label as goods includes some service sectors.

Preferences take the following form:

$$U(c_G, c_S) = \left[a_G c_G^{\frac{\varepsilon-1}{\varepsilon}} + (1 - a_G) (c_S + \bar{c}_S)^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}$$

where c_G and c_S are consumption of goods and services, $0 < a_G < 1$, $\bar{c}_S \geq 0$ and $\varepsilon > 0$. Note that if $\bar{c}_S > 0$, preferences are non-homothetic and, holding prices constant, the expenditure share on services will be increasing in income.²¹ This non-homotheticity is motivated by the cross-sectional analysis in the previous section. Note that households are assumed to not value leisure, since our focus will be on the relative prices of labor given observed supplies.

²⁰We later carry out an exercise to assess how changes in net trade flows by sector affect our key findings.

²¹This is a simple and common way to create differential income effects across the two consumption categories. One can also generate non-homothetic demands in other ways. For example, Hall and Jones (2007) generate an income elasticity for medical spending that exceeds unity through the implied demand for longevity. Boppart (2014), Swiecki (2017) and Comin et al. (2015) all consider more general preferences with the common feature being that income effects associated with non-homotheticities do not vanish asymptotically. This property is likely to be relevant when considering a sample with countries at very different stages of development. Because we focus on a sample of predominantly rich countries, we have chosen to work with the simpler preference structure in order to facilitate transparency of the economic forces at work.

Each of the two production sectors has a constant returns to scale CES production function that uses low- and high-skill labor as inputs:

$$Y_j = A_j \left[\alpha_j H_j^{\frac{\rho-1}{\rho}} + (1 - \alpha_j) L_j^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}} \quad j = G, S$$

where L_j and H_j are inputs of low- and high-skill labor in sector j , respectively, α_j captures skill-biased technical change in sector j and A_j reflects skill-neutral technical change in sector j . Our benchmark specification assumes that ρ , the elasticity of substitution between low- and high-skill labor, is the same in both sectors.²²

Our representation of technical change merits some discussion. Technical change in each sector is two-dimensional, and can be represented in many equivalent ways. Our chosen representation is particularly convenient for the effects that we will emphasize. Specifically, in settings with an aggregate production function the skill premium is affected only to the extent that technical change affects the relative marginal products, and is independent of what happens to overall output. Our analysis will focus on the effects associated with the component of technical change that affects output without affecting relative marginal products. The above representation is convenient relative to common alternatives because changes in α will have a first-order effect on relative marginal products but a dampened effect on output due to the opposing effects embedded in the specification.²³

We emphasize that our representation does not imply that we view changes in the A_j and the α_j as two independent processes; rather, our representation is simply a decomposition of the process of technical change into two components. Any pattern of factor augmenting technical change can be decomposed into these two pieces.

Before proceeding to analyze the equilibrium for our model we comment on the significance of abstracting from capital and trade. By excluding capital we implic-

²²We consider the effects of cross-sectional variation in this parameter in the Online Appendix.

²³More generally, consider a two factor CRS production function $F(H, L)$. One natural representation of technical change is $F(A_H H, A_L L)$. In this case changes in either of the A_i generate first order effects on both output and relative marginal products. Two alternative representations that partially address this are to write either $A_H F(H, \frac{A_L}{A_H} L)$ or $A_L F(\frac{A_H}{A_L} H, L)$. But in each case there is still a first-order effect of changes in A_H/A_L on output. While our specification dampens the effect of skill-biased technical change on output, it does not completely eliminate this effect. We have carried out a sensitivity exercise in which we allocate part of the change in the A_j to the skill-biased component, so that both the direct and indirect (i.e., general equilibrium) effect of skill-biased technical change on aggregate output is exactly equal to zero. This has a modest effect on our results and so is included in the Online Appendix.

itly adopt a somewhat reduced form view of skill-biased technological change. For example, changes in relative demand for skilled labor due to capital-skill complementarity and changes in the price of equipment (as in Krusell et al., 2000) will show up in our model as skill-biased technological change. While it is obviously of interest to understand the underlying mechanics of skill-biased technological change, we believe our results are strengthened by adopting a more expansive notion of skill-biased technological change rather than focusing on a particular mechanism.²⁴

Although our benchmark analysis abstracts from trade, it implicitly captures some potential effects of trade. In particular, changing patterns of trade may affect the composition of production within the low-skill intensive sector due to specialization. If this involves specialization in higher skill sectors within our low-skill intensive sector, our analysis will capture this as skill-biased technical change within the low-skill intensive sector. That is, part of what we measure as skill-biased structural change within the low-skill intensive sector may reflect the effects of trade.

A separate issue is that as trade in services has increased over time, it may also contribute to the changing composition of US production across the low- and high-skill intensive sectors. In Section 7.2 we carry out an exercise to assess the importance of this effect.

3.2 Equilibrium

We focus on a competitive equilibrium for the above economy. The competitive equilibrium will feature four markets: two factor markets (low- and high-skill labor) and two output markets (goods and services), with prices denoted as w_L , w_H , p_G and p_S . We will later normalize the price of low-skill labor to unity so that the price of high-skill labor will also represent the skill premium.

The definition of competitive equilibrium for this model is completely standard, so we move directly to characterizing it. Individuals of skill $i = L, H$ solve

$$\max_{\{c_{Gi}, c_{Si}\}} \left[a_G c_{Gi}^{\frac{\epsilon-1}{\epsilon}} + (1 - a_G) (c_{Si} + \bar{c}_S)^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}}$$

subject to

$$p_G c_{Gi} + p_S c_{Si} = w_i. \tag{1}$$

²⁴Acemoglu and Guerrieri (2008) emphasize that capital accumulation may also be a cause of structural change. In our framework these effects will be captured by changes in the A_j .

Using the first-order conditions of this problem and normalizing w_L to unity, the aggregate expenditure share for services, denoted by e_S satisfies:

$$e_S = \frac{p_S [(1-f)c_{SL} + fc_{SH}]}{1-f + fw_H} = \frac{1}{\left(\frac{1-a_G}{a_G}\right)^\epsilon + \left(\frac{p_G}{p_S}\right)^{1-\epsilon}} \left[\left(\frac{1-a_G}{a_G}\right)^\epsilon - \frac{p_S \bar{c}_S \left(\frac{p_G}{p_S}\right)^{1-\epsilon}}{1-f + fw_H} \right]. \quad (2)$$

This expression illustrates the two forces driving structural change from the perspective of the household: relative prices and income. Specifically, if $\epsilon < 1$, the expenditure share of services increases as p_G/p_S declines, and if $\bar{c}_S > 0$, the expenditure share of services increases as income measured in units of services (i.e., $(1-f + fw_H)/p_S$) increases.

The problem of the firm in sector $j = G, S$ is

$$\max_{\{H_j, L_j\}} p_j A_j \left[\alpha_j H_j^{\frac{\rho-1}{\rho}} + (1-\alpha_j) L_j^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}} - w_H H_j - L_j.$$

Cost minimization plus the requirement that profits be zero in a competitive equilibrium for a firm with a constant returns to scale production function imply an equation for the price of sector j output in terms of the skill premium w_H :

$$\hat{p}_j(w_H) = \frac{1}{A_j} \left[\frac{\alpha_j^\rho}{w_H^{\rho-1}} + (1-\alpha_j)^\rho \right]^{\frac{1}{1-\rho}}. \quad (3)$$

It follows that finding equilibrium prices can be reduced to a single dimension: if we know the equilibrium value of w_H then all of the remaining equilibrium prices can be determined.

Equilibrium requires that all four markets clear: the two markets for output and the two markets for labor. Here we derive an expression for the market-clearing condition for high-skilled labor that contains the single price w_H . Using $H_j/L_j = \left(\frac{\alpha_j}{1-\alpha_j} \frac{1}{w_H}\right)^\rho$, the production function of sector j , and (3), we obtain a sector-specific demand function for high-skilled labor:

$$H_j = \left[\frac{\alpha_j \hat{p}_j(w_H) A_j}{w_H} \right]^\rho \frac{Y_j}{A_j}, \quad (4)$$

which, together with equilibrium in the goods market, yields the market-clearing

condition for high-skilled labor solely as a function of w_H :

$$\begin{aligned} & \left[\frac{\alpha_S \hat{p}_S(w_H) A_S}{w_H} \right]^\rho \frac{f \hat{c}_{SH}(w_H) + (1-f) \hat{c}_{SL}(w_H)}{A_S} \\ & + \left[\frac{\alpha_G \hat{p}_G(w_H) A_G}{w_H} \right]^\rho \frac{f \hat{c}_{GH}(w_H) + (1-f) \hat{c}_{GL}(w_H)}{A_G} = f. \end{aligned} \quad (5)$$

Here we have used $\hat{c}_{ji}(w_H)$ to denote the demand for output of sector j by a household of skill level i when the high-skilled wage is w_H and prices are given by the functions $\hat{p}_j(w_H)$ defined in (3).

3.3 Structural Change and the Skill Premium

In this subsection we derive an analytic expression that summarizes how the sector-specific skill-neutral component of technical change affects the skill premium. To do this we focus on the special case in which preferences are homothetic (i.e., $\bar{c}_S = 0$) and solve for a linear approximation of the model (i.e., equations (1), (2), (3), and (5)). The resulting log-linear expression is:

$$\frac{dw_H}{w_H} = \frac{1-f+fw_H}{(1-f)\tilde{\rho}} (h_S - e_S) (1-\varepsilon) \left(\frac{dA_G}{A_G} - \frac{dA_S}{A_S} \right)$$

where $h_j = H_j/f$, $e_S = p_S [(1-f)c_{SL} + fc_{SH}] / (1-f+fw_H)$, and $\tilde{\rho}$ is given by:

$$\tilde{\rho} = \rho \frac{(1-f+fw_H)}{(1-f)} \left[\left((1-\theta_S) h_S + (1-\theta_G) h_G \right) + \varepsilon (\theta_S - \theta_G) (h_S - e_S) \right],$$

where $\theta_j = w_H H_j / (p_j Y_j)$.

Several results follow. First, proportional changes in the A_j have no impact on the skill premium. (Recall that this derivation assumed homothetic preferences.) Second, changes in the relative value of the A_j will have no effect on the skill premium if $\varepsilon = 1$. (Notably, these results parallel standard results in the structural change literature regarding conditions under which technical change generates structural change.) Third, assuming changes in the relative values of the A_j , and that $\varepsilon \neq 1$, there will be an effect on the skill premium if and only if there is heterogeneity in skill intensity, i.e., if and only if the share of high-skill labor in services, h_S , differs from the expenditure share of services, e_S . In particular, if $\frac{dA_G}{A_G} - \frac{dA_S}{A_S} > 0$, and $\varepsilon < 1$, then the skill premium

will increase if and only if the service sector is more skill intensive, i.e., $h_S > e_S$.

The effect of changes in the A_j on the skill premium are intimately related to structural change: the change in relative prices in the log linearized model is proportional to $\frac{dA_G}{A_G} - \frac{dA_S}{A_S}$ and the extent to which this change in relative prices affects expenditure shares is dictated by the value of $(1 - \varepsilon)$.

For future purposes it is also of interest to derive an expression for the effect of a change in the supply of skill on the skill premium. Our log linearization yields:

$$d \log w_H = - \left(\frac{1}{\tilde{\rho}} \right) d \log \left(\frac{f}{1-f} \right) \quad (6)$$

where $\tilde{\rho}$ is as defined above.

Equation (6) highlights the extent to which our two-sector model generalizes the expression for the elasticity of the skill premium to a change in the supply of skills relative to a one sector model. In a one sector model this elasticity is completely determined by the elasticity of substitution in production and equals $-1/\rho$. In our two sector model, the effective aggregate elasticity of substitution between the two types of labor is potentially different due to the fact that one can substitute labor across sectors. This can either amplify or dampen the effective elasticity of substitution relative to a one sector model.²⁵ If $\varepsilon = 0$, the two-sector elasticity is smaller than the elasticity of substitution in production, $\tilde{\rho} < \rho$, but for ε sufficiently high the reverse holds, i.e., $\tilde{\rho} > \rho$. We will use this expression in the next section when we calibrate the value of ρ .

4 Calibration

In this section we calibrate the model of the previous section so as to be consistent with the salient features of structural change, growth, and the changes in the skill premium under the assumption that the driving forces are changes in technology (both skill-biased and skill-neutral) and changes in the relative supply of skill.²⁶ In particular,

²⁵Intuitively, we see that $\tilde{\rho} = \rho$ if we ignore the general equilibrium impacts of the skill premium on the price of sectoral output and of the changes in the supply of high-skill labor on the demand for high-skill labor. That is, if we set $\rho(\theta_S h_S + \theta_G h_G) + \varepsilon(\theta_S - \theta_G)(h_S - c_S) = 0$ and $f(w_H - 1) = 0$, respectively. Of course, $\tilde{\rho} = \rho$ would also hold if these two general equilibrium effects happen to perfectly offset each other.

²⁶To the extent that factors such as changes in the minimum wage and unionization affect the skill premium, our analysis will identify them as changes in skill-biased technical change. This was also

we will use the above model to account for observed outcomes at two different points in time, that we denote as 0 and T for the initial and terminal periods respectively. Consistent with the existing literature on technological change and the skill premium, we do not allow the parameter ρ to change over time. We also assume that preferences are constant over time.

Calibrating the model in the initial and terminal period requires assigning values for 14 parameters. Nine of these are technology parameters: four values of the α_j (two in each period), four values of the A_j (two in each period), and ρ . Three are preference parameters: ε , a_G and \bar{c}_S . Lastly we have the value of f at the initial and terminal dates. The two initial values of the A_j represent a choice of units and the initial and final values of f will be measured directly from the data. We will calibrate the elasticity parameter ρ in accordance with existing estimates, appropriately filtered through our model, as summarized by equation (6). (We describe this in more detail at the end of the section.) This leaves nine parameters to be calibrated, six technology parameters and three preference parameters.

Our calibration procedure will proceed in two steps. The first step describes how we determine the six technology parameters independently of the three preference parameters. Having determined the six technology parameters we then describe how we determine the three preference parameters.

4.1 Calibrating Technology Parameters

In this section we show that the six technology parameters can be determined independently of the three preference parameters if we target the following eight values from the data: the initial and final values for factor shares in both sectors, the initial and final value added shares for the two sectors, the initial and final value of the skill premium, the change in the relative price of the two sectors, and the overall growth rate of the economy.²⁷

To measure these targets in the data we rely on the World KLEMS data for the

the case for the analysis of Katz and Murphy (1992). Our estimate of the contribution of skill-biased technical change should be understood as including the effects of these other factors.

²⁷Recall that in Section 2 we reported evidence for changes in both sectoral value added shares as well as sectoral compensation shares. Although both displayed the same pattern, the changes in compensation shares were somewhat larger. Our baseline calibration uses data on value added shares, but in Section 3 of the Online Appendix we repeat our main exercise using compensation shares instead. Our main message is robust to this change, and intuitively, implies a larger role for skill biased structural change.

U.S. for the years 1977 and 2005.²⁸ This period is of particular interest, since 1977 effectively marks a local minimum in the skill premium (see Acemoglu and Autor, 2011, for earlier data), after which it secularly increases.²⁹

Many of the targets have obvious counterparts in the data and so require no discussion. But two issues merit some discussion. The first concerns the fact that because our model does not include investment, we implicitly assume that output value-added shares reflect consumption value-added shares. We show in our sensitivity analysis that adjusting the data as in Herrendorf et al. (2013) to compute consumption value-added shares has virtually no impact on the targets used for calibration.

The second issue concerns the targets for the labor variables. As we discuss in more detail later on, our procedure for decomposing compensation into price and quantity is an important point of departure from the analysis in Katz and Murphy (1992), so we next provide some detail on our method.

World KLEMS contains data on labor compensation per hour worked and average hours worked by week by industry, educational attainment, class, gender, and age groupings, as well as the number of employed individuals in each of these groupings.³⁰ Consistent with our calculations in Section 2, we combine all workers with less than college completion into our classification of low-skilled, and all workers with college completion or more into our classification of high-skilled to calculate labor income

²⁸We use World KLEMS rather than EUKLEMS in this exercise to facilitate comparison with the work by Katz and Murphy (1992). They base their analysis on the CPS, and it turns out that the micro data underlying World KLEMS is much closer to the CPS data than the micro data underlying EUKLEMS. The reason for the difference is that EUKLEMS makes adjustments so as to make their data match data from the BEA. We note however, that although the two datasets provide slightly different answers for the shift-share calculations of Katz and Murphy, our model based results are effectively unchanged if we instead use the EUKLEMS data to provide all of our calibration targets. We continue to use EUKLEMS for cross-country comparisons because this cannot be done in World KLEMS.

²⁹We choose 2005 as our terminal date because this is the last period available consistently across datasets.

³⁰Until 1992 educational attainment is based on years of schooling and classified into six categories “less than high school”, “some high school”, “high school graduates”, “some college”, “college graduates”, and “more than college graduates”. After 1992, this classification changes to highest level achieved, being the categories “8th grade or less”, “grades 9-12 no diploma”, “high school graduates”, “some college no degree, associate degree”, “BA,BS”, and “more than BA”. Whenever we need to compute consistent time series using these categories, we perform the adjustment suggested in Jaeger (1997).

There are two classes of workers (employees and self employed) and eight age groups (14-15, 16-17, 18-24, 25-34, 35-44, 45-54, 55-64, and 65 and over). Weekly hours are normalized so that weeks worked per year total 52.

shares by skill at both the aggregate and sectoral level. We use the same sectoral classification as in Section 2.

Setting targets for the skill premium and the relative supply of skilled workers requires that we decompose labor payments into price and quantity components. If all workers within each skill type were identical then we could simply use total hours as our measure of quantity, but given the large differences in hourly wage rates among subgroups in each skill type this seems ill-advised. Instead, we assume that each subgroup within a skill type offers a different amount of efficiency units per hour of work.³¹ We normalize efficiency units within each skill type by assuming one hour supplied by a high school-educated prime-aged (i.e., aged 35-44) male is equal to one efficiency unit of low-skill labor and that one hour supplied by a college-educated prime-aged (i.e., aged 35-44) male is equal to one efficiency unit of high-skill labor.³² With this choice of units, the skill premium is defined as the ratio of college-educated to high school-educated prime-aged (i.e., aged 35-44) male wages. This premium rises from 1.33 in 1977 to 1.88 in 2005.³³ Note that our implicit assumption is that differences in wages between different demographic groups within a given skill category reflect differences in efficiency units. This interpretation is consistent with standard practice in the literature on heterogeneous agent models.

We infer f using the identity that the ratio of labor compensation equals the product of the skill premium and the relative quantity of high- to low-skill labor (f and $1 - f$, respectively).³⁴ This procedure implies that high-skill labor was 21% of total labor supply in 1977 and rose to 32% in 2005.

Table 2 summarizes the values that will be used to calibrate the technology parameters.

³¹The Online Appendix contains more details on our procedure.

³²While one could obviously normalize units by choosing other reference groups, this group seems most natural since its uniformly high rate of participation over time minimizes issues associated with selection.

³³Comparing earnings of full time workers using CPS data, Figure 1 in Acemoglu and Autor (2011) indicates values of 1.48 and 1.89 for 1977 and 2005 respectively. Our measure indicates a fourteen percentage point greater increase. This difference basically reflects the fact that Acemoglu and Autor compute a fixed-weight, composition-adjusted average wage for high school and college graduates of different experience, race, and gender groups. If we redo their analysis with CPS data but using only male workers aged 35-44, we find a 52 percentage point increase in the skill premium, consistent with our measured increase using World KLEMS data.

³⁴Equivalently, one could compute efficiency units of each skill type by using relative wages within each skill group to infer efficiency units and directly aggregating efficiency units.

Table 2

Values Used to Calibrate Technology Parameters

f_0	f_T	w_{H0}	w_{HT}	$\% \Delta \frac{p_S}{p_G}$	$\% \Delta Y$	θ_{G0}	θ_{GT}	θ_{S0}	θ_{ST}	$\frac{P_{S0}C_{S0}}{Y_0}$	$\frac{P_{ST}C_{ST}}{Y_T}$
0.21	0.32	1.33	1.88	45.6	80.8	0.18	0.36	0.50	0.65	0.25	0.39

We now describe the details of how these values are used to determine the values of the six technology parameters. We begin with the determination of the α_{jt} . Given a value for ρ , the four values of the α_{jt} are pinned down by sectoral factor income shares and the skill premium, w_{Ht} . To see this, from equations (3) and (4) note that the share of sector j income going to high-skill labor, $\theta_{jt} = \frac{w_{Ht}H_{jt}}{\hat{p}_j(w_{Ht})Y_{jt}}$, is

$$\theta_{jt} = \frac{\alpha_{jt}^\rho}{\alpha_{jt}^\rho + (1 - \alpha_{jt})^\rho w_{Ht}^{\rho-1}}$$

Therefore, given ρ , the skill premium w_{Ht} , and data for θ_{Hjt} , the value of the α_{jt} are given by:

$$\alpha_{jt} = \frac{1}{1 + \frac{1}{w_{Ht}^{(\rho-1)/\rho}} \left(\frac{1 - \theta_{Hjt}}{\theta_{Hjt}} \right)^{\frac{1}{\rho}}}.$$

Next we determine the values of the A_{jt} 's. As noted previously, the two values in period 0 basically reflect a choice of units and so can be normalized. We will normalize A_{S0} to equal one, and given the calibrated values for the α_{j0} and the value of w_{H0} , we choose A_{G0} so as to imply $p_{G0}/p_{S0} = 1$. A convenient implication of this normalization is that p_{GT}/p_{ST} is not only the level of relative price in period T but is also the change in the relative price between periods 0 and T .³⁵

As is well known in the literature, with identical Cobb-Douglas sectoral technologies, relative sectoral prices are simply the inverse of relative sectoral TFPs, so the change in relative prices would therefore determine the values of the two A_{jT} 's up to a scale factor.³⁶ This precise result does not apply to our setting because of sectoral heterogeneity in the α_{jt} 's. Nonetheless, there is still a close connection between relative sectoral prices and relative values of the A_{jt} . In particular, using equation (3)

³⁵While our main results will only use information from the initial and final periods, we note that the procedure described here can be used to uncover the entire sequence of technology parameters from period 0 to period T .

³⁶This same relation holds more generally, and in particular would also apply if the sectoral production functions are CES with identical parameters.

for the two sectors we have:

$$\frac{A_{Gt}}{A_{St}} = \frac{p_{St}}{p_{Gt}} \left[\frac{\alpha_{Gt}^\rho}{w_{Ht}^{\rho-1}} + (1 - \alpha_{Gt})^\rho \right]^{1/(1-\rho)} \cdot \left[\frac{\alpha_{St}^\rho}{w_{Ht}^{\rho-1}} + (1 - \alpha_{St})^\rho \right]^{-1/(1-\rho)}. \quad (7)$$

The scale factor influences the overall growth rate of the economy between periods 0 and T , so we choose this scale factor to target the aggregate growth rate of output per worker. Note that to compute aggregate output at a point in time (and thus also the growth rate in aggregate output) it is necessary to know the sectoral distribution of output. The relations imposed thus far guarantee that maximum profits will be zero in each sector, but they do not determine the scale of operation. Intuitively, the split of activity across sectors at given prices will be determined by the relative demands of households for the two outputs. Below we describe how preference parameters are chosen to match the sectoral distribution of value added at both the initial and final date. At this stage we simply assume this split is the same as in the data.

To this point we have identified all of the technology parameters conditional on a value of ρ . We postpone a detailed discussion of the calibration of ρ until the end of this section, but note here that for our benchmark analysis we set $\rho = 1.53$ and that filtered through equation (6), our value implies an effective aggregate elasticity very close to the one used by Katz and Murphy (1992). Table 3 presents the benchmark calibrated values for the technology parameters.

Table 3

Calibrated Technology Parameters ($\rho = 1.53$)					
α_{G0}	α_{S0}	α_{GT}	α_{ST}	A_{ST}/A_{S0}	A_{GT}/A_{G0}
0.29	0.53	0.46	0.65	1.45	2.35

We note three features from this table. First, and not surprising given the way in which we grouped industries into the two sectors, the weight on low-skill labor is greater in the goods sector than in the service sector at both dates. Second and more interesting is that in both sectors technological change has an important component that is skill-biased. In fact, the level rise in α is greater for the goods sector than the service sector (0.17 versus 0.12). And third, neutral technological progress is much greater in the goods sector than in the service sector. The average annual growth

rate of A_{Gt} is 2.99%, while the average annual growth rate of A_{St} is only 1.29% per year.

4.2 Calibrating Preference Parameters

We now turn to the issue of determining values for the three preference parameters: a_G , \bar{c}_S and ε . While the previous subsection showed that technological change can be inferred without specifying any of the preference parameters, we cannot evaluate some of the counterfactual exercises of interest without knowing how relative demands for the sectoral outputs are affected by changes in prices.

The calibration of the A_{jT} used information about sectoral expenditure shares without guaranteeing that observed expenditure shares were consistent with household demands given prices. Requiring that the aggregate expenditure share for goods and services are consistent with the observed values in the data for the initial and terminal date provides two restrictions on the three preference parameters. Loosely speaking, given a value for a_G , requiring the model to match the initial and final value-added shares requires that the model match the overall amount of structural change but does not determine the relative contribution of income effects and relative price effects. These are in turn dictated by the values of \bar{c}_S and ε . Knowledge about one of these parameters would allow us to infer the other.

Earlier in this paper we presented evidence on the effect of income on the relative expenditure share. We also emphasized that estimates based on the CEX should be treated with caution given they do not include government expenditure and that this is an important component of overall spending on skill-intensive sectors such as health and education.

Alternatively, the empirical literature has provided estimates of ε that correspond to the categories of “true” goods and “true” services, but not for our definitions of the two sectors that are based purely on skill intensity. However, given that our goods sector does contain all of the industries that produce goods, while our service sector does consist entirely of service sector industries, information about the elasticity of substitution between “true” goods and “true” services is plausibly informative about the empirically plausible range of values for ε in our model. Recalling that the objects in our utility function reflect the value added components of sectoral output, the relevant estimates in the literature would include Herrendorf et al. (2013), Buera and

Kaboski (2009), and Swiecki (2017). All of these studies suggest very low degrees of substitutability.³⁷ Based on these studies we think a reasonable range of values for ε is 0 to 0.50.

In light of the above partial information about income and substitution effects, we proceed as follows. We consider three values of ε from the above range: 0.01, 0.10, and 0.50. For each of these values for ε , we use equation (2) to determine values for a_G and \bar{c}_S by requiring the model to match the initial and final sectoral value added shares. Given these values we compute the implied income elasticity of the relative expenditure share for the skill-intensive sector by comparing the consumption expenditure shares and incomes of low- and high-skill workers in our model. When $\varepsilon = 0.10$, the implied income elasticity for the skill-intensive expenditure share is 0.085, which is close to the value in column 4 of Table 1. We choose this as our benchmark specification.

Our main results turn out to be quite insensitive to the relative importance of income and substitution effects in generating the observed amount of structural change. To allow us to explore this sensitivity even more fully we will also consider the case of $\varepsilon = 1.00$, which implies that structural change is entirely caused by income effects, since relative prices have no impact on expenditure shares when $\varepsilon = 1.00$ if preferences are homothetic. Table 4 shows the calibrated preference values for each of the values of ε .

Table 4

Calibrated Preference Parameters			
	ε	a_G	\bar{c}_S
Benchmark	0.10	0.99	0.11
Low ε	0.01	1.00	0.09
Intermediate ε	0.50	0.49	0.24
High ε	1.00	0.26	0.87

The qualitative patterns in this table are intuitive. Recall that the calibration procedure implies that in each specification the changes in income, relative prices and aggregate expenditure shares are the same. Consider the changes as we move from $\varepsilon = 0.10$ to $\varepsilon = 0.50$. This increases the elasticity of substitution between the

³⁷Comin et al. (2015) redo the exercise in Herrendorf et al. (2013) for a more general class of preferences and find an elasticity of substitution that is somewhat higher, around 0.50, which is our intermediate case.

two goods, implying a smaller response in relative expenditure shares. To compensate for this smaller effect, the impact of income changes on relative expenditure shares must increase, implying a higher value for \bar{c}_S . The higher value for \bar{c}_S will in turn lead to a lower expenditure share on services in the initial period, thereby requiring a lower weight, a_G , on the consumption of goods.

Consistent with the literature that considers “true” goods and “true” services, we also find that some income effects are needed to rationalize the data, as \bar{c}_S remains positive even when $\varepsilon = 0.01$, which is effectively the case of Leontief preferences and serves to maximize the role of relative price effects.³⁸

4.3 Calibrating ρ

We are now in a position to describe our procedure for calibrating the value of ρ . Our procedure follows closely the one originally adopted by Katz and Murphy (1992), and followed by many authors subsequently. They assumed an aggregate CES production function:

$$Y_t = A_t[\alpha_t H_t^{1-\frac{1}{\rho}} + (1 - \alpha_t)L_t^{1-\frac{1}{\rho}}]^{\frac{\rho}{\rho-1}}$$

This specification implies:

$$\Delta \log \left(\frac{w_{Ht}}{w_{Lt}} \right) = \Delta \log \frac{\alpha_t}{(1 - \alpha_t)} - \left(\frac{1}{\rho} \right) \Delta \log \left(\frac{H_t}{L_t} \right)$$

Their strategy was to assume that technical change proceeded at a constant rate and to identify ρ from the following regression:

$$\log \left(\frac{w_{Ht}}{w_{Lt}} \right) = a_1 + a_2 t + a_3 \log \left(\frac{H_t}{L_t} \right) + \varepsilon_t. \quad (8)$$

They concluded that $\rho = 1.41$.

We follow this same strategy using our data and time period, but note two differences in implementation. First, as noted earlier, we use a different procedure to measure efficiency units of labor. Second, recalling equation (6), the above regression

³⁸To have better sense of magnitudes, in the benchmark case the value of the non-homotheticity parameter relative to GDP $p_S \bar{c}_S / (1 - f_H + f_H w_H) = 0.24$ and 0.17 in the initial and final periods, respectively.

will identify the value of $\tilde{\rho}$, which we then use to infer the value of ρ .³⁹ Note that the mapping from $\tilde{\rho}$ to ρ depends on the preference parameter ε , so that our procedure implies a different value of ρ for each profile of preference parameters.

Despite the difference in time periods and measurement of labor, our estimated value of a_3 in equation (8) is very close to that obtained by Katz and Murphy (1992); they obtained a value of -0.709 while we obtain -0.708 . For our benchmark specification with $\varepsilon = 0.10$ the implied value of ρ is 1.53, which is also quite close to the value of 1.41 used by Katz and Murphy (1992). For ε in the range of 0 to 0.50, the variation in ρ is not that large, varying from 1.42 to 1.55. When $\varepsilon = 1.0$ the change is more substantial, as the implied value of ρ is 1.18 for this case.

5 Results

The procedure described in the previous section implies that our calibrated model will perfectly account for the observed change in the skill premium between 1977 and 2005. In this section we use the calibrated model to decompose changes in the skill premium into parts due to the exogenous driving forces in the model. Our primary objective is to decompose the effect of technical change on the skill premium into a piece due to the skill-biased component of technological change and a second piece due to the sector-specific skill-neutral component of technological change.

5.1 Sources of Change in the Skill Premium

We begin by decomposing the change in the skill premium between 1977 and 2005 into three pieces. The first piece represents the effect of changes in the supply of skill (i.e., changes in f). Consistent with equation (6), an increase in the supply of skill holding technology constant will lead to a decline in the skill premium. The value of f increases from 0.21 in 1977 to 0.32 in 2005. In our benchmark specification (i.e., $\varepsilon = 0.10$) this increase in f holding all else constant would have led to a reduction of w_H from 1.33 to 0.87.⁴⁰

In reality (and in our calibrated model) the skill premium increased from 1.33 in

³⁹For this step we solve for the implied value of ρ at the initial period in our sample. The implied value of ρ is roughly the same if we take the final period instead.

⁴⁰Because we calibrate ρ so as to hold the effective local elasticity of substitution constant as ε varies, this change is very nearly identical for all of our profiles for preference parameters.

1977 to 1.88 in 2005. As just noted, if the only change had been a change in the supply of skill, then w_H would have decreased to 0.87 in 2005. From this we conclude that technical change generated an increase in the demand for skill that collectively increased the skill premium from 0.87 to 1.88, an increase of 1.01. Our next goal is to decompose this change of 1.01 into one part due to the skill-biased component of technical change (i.e., changes in the α_j) and one part due to the sector-specific skill-neutral change component of technical change (i.e., changes in the A_j). In the next subsection we show that the changes in the A_j are intimately related to structural change.

A natural way to evaluate each of these effects is to start from the economy with initial technology settings and $f = f_T$, which would imply $w_H = 0.87$, and evaluate the effect on w_H of moving to final values for either the α_j or the A_j . If the model were linear, these two effects would necessarily add up to the total effect of changes in technology. However, it turns out that there is a relatively small positive residual that reflects interactions between the two different changes in technology. Table 5 presents the results of this exercise.

Table 5
Effect of Technological Change on the Skill Premium
US, 1977-2005

	$\varepsilon = 0.01$	$\varepsilon = 0.10$	$\varepsilon = 0.50$	$\varepsilon = 1.00$
(i) Total Δw_H due to Technology	1.00	1.01	1.01	1.02
(ii) Δw_H due purely to the A_j	0.18	0.18	0.17	0.15
(iii) Δw_H due purely to the α_j	0.77	0.77	0.76	0.76
(iv) Δw_H due to interaction	0.05	0.06	0.08	0.09
(v) % Contribution of the A_j	18.1 – 23.8	18.0 – 24.0	17.2 – 24.5	15.3 – 24.4

The first row of Table 5 presents the total increase in the skill premium due to all sources of technical change and represents the difference between the actual skill premium in 2005 and our model implied counterfactual for what the skill premium would have been if the only change had been the supply of skill. As noted earlier, this is effectively the same for all of the specifications.

The next three rows in Table 5 present the size of the change due purely to changes in the A_j , due purely to changes in the α_j , and the interaction term. The final row

shows the range for the percent contributions of the change in the A_j as we vary the fraction of the interaction term allocated to changes in A_j from zero to one.

Focusing on the benchmark calibration ($\varepsilon = 0.10$) case for now, the final row of Table 5 shows that the change in the A_j account for between 18 and 24 percent of the overall change in the skill premium due to technological change. Put somewhat differently, according to our calibrated model, if skill-biased technical change had been the only force affecting the relative demand for skill then the skill premium would have increased by only 31 – 37 percentage points over the period 1977 to 2005 instead of increasing by 55 percentage points.

While our results imply that the skill-biased component of technical change is the dominant source of changes in the skill premium, some care should be exercised in interpreting this. As we noted earlier, the literature has noted that changes due to factors such as minimum wages, unionization and trade will be included as reflecting skill-biased technical change. Some estimates claim that as much as half of measured skill-biased technical change might be due to these other factors. (See, for example DiNardo et al., 1996) With this in mind, our results suggest a much less dominant role for skill-biased technical change.

Looking at the other columns in Table 5 we see that the basic message about the relative importance of the change in the A_j is fairly similar across the specifications, with the overall range being 15.3 – 24.5 percent. This is reassuring given the lack of definitive values for ε and \bar{c}_S . However, the table does reveal a tendency for the pure effect of the A_j to decrease as we increase the relative importance of income effects as a driver of structural change. Specifically, as ε increases from 0.01 to 1.00 the pure effect of the A_j decreases from 18.1 to 15.3. It is of interest to understand why the effect is relatively small. To be sure, this result partly reflects our calibration strategy: in all of our specifications the overall change in income and relative prices is the same as in the data, and in each case the values of ε and \bar{c}_S are such that we necessarily account for the same change in expenditure shares as in the data. But as we detail in Appendix 1, the relatively small effect of changes in ε and \bar{c}_S also reflects properties of the data that do not rely on these details of the calibration procedure. For example, it is relevant that the change in the α_j are found to have virtually zero effect on relative prices, but this result reflects the fact that two opposing effects are roughly offsetting in the data. On the one hand, because the service sector is more intensive in high skilled labor, a given change in the α_j will have a larger impact

on service sector productivity. But we also find that the change in α_G exceeds the change in α_S , and roughly offsets this first effect.

We conclude that our message of the importance of changes in the A_j for changes in the skill premium is robust to the relative importance of income and substitution effects in generating structural change, and even our estimates are fairly robust; that is, conditional on our model being calibrated so as to generate the amount of structural change that we see in the data, it is less important to determine the mix of income and substitution effects that leads to this change.

5.2 Sources of Structural Change

In the introduction we stressed the fact that aggregate production function analyses abstract from compositional changes, and that a key objective of our analysis was to assess the quantitative importance of the compositional changes that are associated with the process of structural transformation during development. The previous calculations decomposed the overall changes in the skill premium due to technology into parts due to the α_j and the A_j . In order to make the connection between this decomposition and compositional changes it is necessary to examine the connection between structural change and the components of technical change.

To do this we start by assessing the amount of structural change that would have occurred had there not been any change in the A_j . Results are shown in Table 6.

Table 6
Value Added Share of the Skill-Intensive Sector
US, 1977-2005

	$\varepsilon = 0.01$	$\varepsilon = 0.10$	$\varepsilon = 0.50$	$\varepsilon = 1.00$
Model 1977	0.25	0.25	0.25	0.25
Model 2005	0.39	0.39	0.39	0.39
Model 2005 with fixed A_j	0.21	0.21	0.22	0.24

The first two rows of the table remind us that the skill-intensive sector grew significantly between 1977 and 2005, increasing its share of value added from 25 percent to 39 percent. Recall that our calibrated model perfectly replicates the change in the data. The third row shows what would have happened if there had not been any changes in the A_j . Significantly, this would have led to a decline in the value-added share for the skill-intensive sector. Given that (more than) all of the observed

structural change is due to changes in the A_j , we think it is appropriate to identify the channel through which the A_j affect the skill premium as reflecting compositional change.

It is significant that the pure effect of the change in the A_j on the value added share of the skill-intensive sector is actually greater than the overall observed change: 0.18 versus 0.14. That is, the amount of structural change generated by the changes in the A_j is more than twenty percent larger than the amount of structural change observed in the data. It follows that the amount of observed structural change in the data is not a good estimate of the amount of structural change induced by the change in the A_j . The significance of this will be highlighted in the next section when we contrast our model-based findings with those found by Katz and Murphy (1992) using shift-share methods.

The third row of Table 6 reported the combined effect of the changes in f and the α_j on the value-added share of the skill-intensive sector. It is also of interest to assess the role of each of these changes separately. It turns out that the changes in the α_j and f contribute equally to this modest decline. If we change f from its 1977 value to its 2005 value but holding technology constant, we see a decline in the value added share of the skill-intensive sector to 0.23. The reason for this decrease is that the increase in the supply of skill lowers the skill premium, thereby reducing income and lowering the relative price of the skill-intensive sector which is more intensive in the use of skilled labor. Both of these effects serve to shift expenditure away from the skill-intensive sector. It is interesting to note that an increase in the supply of high-skill labor does not by itself lead to an expansion of the sector that is more intensive in its use of high-skill labor.

The changes in the α_j also produce a modest decline in the value-added share of the skill-intensive sector. This reflects the net effect of several opposing effects. Because this sector uses high-skill labor more intensively, a uniform increase in α would have a larger productivity effect on it, thereby lowering its relative price and shifting expenditure to the goods sector. But, as noted earlier, the increase in α_G is somewhat larger than the increase in α_S . The increase in the α_j also lead to an increase in the skill premium, which also tends to increase the relative price of the skill-intensive sector.

6 Comparison With the Literature

In the previous section we argued that changes in the composition of demand driven by technical change have played a significant role in the overall increase in the demand for skill. This finding stands in sharp contrast to previous findings in the literature, specifically those in Katz and Murphy (1992) and Leonardi (2015). In this section, we examine the reasons behind these different conclusions.

6.1 Comparison With Katz and Murphy (1992)

We begin by examining how our results compare with those of Katz and Murphy (1992) (hereafter KM). They employ a shift-share method to assess the contribution of changes in sectoral composition to the overall increase in the demand for skill. Specifically, their “Between Industry Demand Shift” for group k measured relative to base year employment of group k in efficiency units ΔX_k^d is given by

$$\Delta X_k^d = \frac{\Delta D_k}{E_k} = \sum_j \left(\frac{E_{j,k}}{E_k} \right) \left(\frac{\Delta E_j}{E_j} \right) = \frac{\sum_j \phi_{j,k} \Delta E_j}{E_k},$$

where E_k is group k 's employment measured in efficiency units and $\phi_{j,k} = (E_{j,k}/E_j)$ is group k 's share of total employment in efficiency units in sector j in the base year. The implied change in the relative demand for skill associated with changes in the sectoral distribution of employment can in turn be used to infer the implied change in the skill premium by using their estimate of the elasticity of substitution between low and high-skilled labor. They conclude that changes in sectoral composition accounted for only 10.6% of the increase in the skill premium between 1979 and 1987.

There are many differences in details between their study and ours: the data sources are different (CPS vs World KLEMS), the measure of payments to workers are different (weekly earnings for full time workers versus compensation per hour), the time periods are different (1963-1987 vs 1977-2005), and the level of aggregation is different (50 sectors vs 2 sectors). In the appendix we report a series of detailed calculations to show that none of these differences is of first-order significance in explaining the very different results. In particular, when we redo the analysis of KM using our data, our time period, and our level of aggregation, we find that changes in sectoral composition account for only 10.3% of the increase in the skill premium.

This result is shown in row (i) of Table 7.

A less apparent difference is that our analysis targeted changes in sectoral composition based on changes in value added shares, whereas the KM procedure measures changes in sectoral composition based on labor compensation shares. This is significant because the change in sectoral compensation shares is greater than the change in sectoral value-added shares.⁴¹

To make the KM numbers directly comparable to ours we redo the KM analysis but using changes in value-added shares. This turns out to have a significant quantitative impact. In particular, row (ii) of Table 7 shows that redoing the KM shift share calculation with value-added shares as sectoral weights reduces their estimated contribution of compositional changes by almost five percentage points, to 5.7%. It is this value that should be compared with our estimated range of 18 – 24%.⁴²

In what follows we show that there are two key differences that account for the fact that our estimate is between 2.5 and 3 times larger than theirs. The first key difference is the method for measuring efficiency units of labor. The second key difference is that our results rely on model based simulation rather than shift share calculations. We discuss each in turn.

Table 7
Comparison With Katz-Murphy (1992)

	Years	Data	Efficiency Units	#Sectors	Method	Weights	Contribution
(i)	1977-05	WK	KM	2	Shift-Share	Wages	10.3%
(ii)	1977-05	WK	KM	2	Shift-share	VA	5.7%
(iii)	1977-05	WK	BKRV	2	Shift-share	VA	12.9%
(iv)	1977-05	WK	BKRV	2	Model-based	VA	18.0 – 24.0%

As noted earlier, within each skill category, we use relative compensation to measure relative efficiency units supplied by individuals and supplied to sectors. Im-

⁴¹We noted this feature of the data in Section 2. For the US, the value added share of the high-skill sector increased from 0.25 to 0.39 between 1977 and 2005 whereas the compensation share of the high-skill sector increased from 0.27 to 0.47.

⁴²Alternatively, we could have redone our benchmark calibration exercise to target the change in compensation shares rather than the change in value added shares. If we do this we find a range of 26.4 – 35.9% for the contribution of changes in the A_j to the change in the skill premium. In both cases our model based approach implies an effect that is between 2.5 and 3 times larger than the corresponding estimate based on the KM shift share calculation. Details of this alternative exercise are included in the Online Appendix.

portantly, we allow for efficiency units to vary across workers within a given education/age/gender cell, since this is how we account for compensation differences within a given cell. In contrast, KM assume that all workers within a given cell supply the same number of efficiency units and measure relative quantities of skilled and less skilled labor without using individual- or sector-specific data on compensation.⁴³ As shown by Row (iii) of Table 7, this turns out to have very significant implications, increasing the estimate based on the KM shift share methodology from 5.7% to 12.9%.⁴⁴

We now turn to the second key difference: our use of a model-based procedure. Row (iv) of Table 7 shows the significance of moving from KM's shift-share analysis to our fully solved general equilibrium evaluation of exogenous shifts in technology parameters. KM acknowledge that their method might underestimate the underlying contribution of demand shifts if other factors, e.g., the rise of the skill premium due to skill-biased technical change, served to dampen the reallocation to skill-intensive sectors.⁴⁵ But they are unable to quantify the extent to which they underestimate the effect. Our model-based method enables us to actually quantify this bias.

Importantly, because we derive endogenous changes in composition as the result of exogenous changes in model primitives, we can map changes in primitives to changes in composition and changes in the skill premium, rather than trying to map changes in composition into changes in the skill premium. To understand the reason that this matters, note that the shift-share calculation uses observed changes in composition to evaluate the effect of compositional changes. But as we showed at the end of the previous section, the change in composition that is associated with changes in the A_j was significantly larger than the observed change in composition. This implies that the shift share calculation will necessarily underestimate the effect of the change in the A_j . As a final remark, we note that our analysis uses a global solution of the model, whereas shift share calculations are inherently based on local approximation.

In summary, while there are many small variations, there are two important fac-

⁴³The Online Appendix provides a framework for thinking about the issue of measuring labor services and details the differences between our procedure and that of Katz and Murphy (1992).

⁴⁴Because our calibrated model replicates all of the values in the data that go into this calculation it follows that this value also reflects what the KM procedure would infer from our model generated data.

⁴⁵Bound and Johnson adjust for the increase in the relative supply of high-skilled labor without accounting for the fact that the relative wage nevertheless rose. This appears to account for their much lower estimate than KM.

tors that explain why we find a substantially larger role for skill-biased structural change in accounting for increases in the skill premium relative to what the earlier literature attributed to industrial composition. The first is that we use wage data to control for unobservable differences among workers within a cell. This implies a larger increase in the demand for efficiency units by the skill intensive sector, thereby increasing the potential impact of compositional changes on the relative demand for skill. The second is that our structural approach allows us to precisely disentangle the role of different driving forces by solving an explicit model-based, globally-solved counterfactual associated with changes in exogenous technology parameters. Each of these factors plays a key role. Measurement differences alone account for a difference of around 7 percentage points, and the use of a model based procedure implies a difference of between 5 and 11 percentage points.

6.2 Comparison With Leonardi (2015)

Leonardi (2015) also asks if changes in composition might be an important mechanism through which some changes in economic primitives lead to changes in the skill premium. Differently than us, he finds that these effects are relatively small. In particular, his exercise finds that the channel of compositional shifts explains approximately 6.5% of the relative demand shift in the US between 1984 and 2002 (see Table 6 in Leonardi, 2015). In this section we discuss the reasons for the apparently different findings and show that there is no inconsistency.

As a first step it is important to summarize the calculations in Leonardi (2015). He specifies a two-sector model very similar to ours. One difference is that in his model, high-skill workers have a relatively higher expenditure share for the output of the skill-intensive sector for two reasons. First, as in our framework, preferences are non-homothetic and the income elasticity for this sector's output is greater than one. Second, he also allows for preferences to differ across low and high-skill workers, and in particular, allows high-skill workers to place greater value on the output of the skill-intensive sector. His calculations are based on a log linear approximation of the demand system generated by this model.

His result comes from the following calculations. First, he calculates the counterfactual percentage change in the skill premium induced by a pure increase in the relative supply of skilled labor. Second, he considers an alternative version of the

model in which preferences of high and low-skill workers are identical and homothetic, and then repeats the previous calculation, i.e., calculates the counterfactual percentage change in the skill premium that would have occurred from a pure increase in the relative supply of skilled workers. His estimate for the effect of demand shifts is calculated by taking the difference in the two percentage changes just calculated. Comparing this to the total percentage change in the skill premium he arrives at 6.5% for the contribution of demand shifts to changes in the skill premium.

Both our paper and Leonardi (2015) present model based calculations about the effect of changes in model primitives on the skill premium that manifest themselves via changes in sector composition. Moreover, the two analyses employ very similar models, with the lone difference being that Leonardi (2015) allows high-skilled workers to have different preferences. But importantly, the two papers focus on different changes in fundamentals. Whereas Leonardi's calculation isolates compositional effects that result from a change in the relative supply of skills, we isolate compositional effects that result from the sector-specific skill-neutral component of technical change. Because the two exercises isolate the effects of different shocks, the different results do not reflect any inconsistency.

To pursue this further we can use our model to carry out the same calculation as Leonardi. The answer that we obtain varies depending upon the profile of preference parameters that we use, and ranges from 0.3% to 5.6%, with the 5.6% value coming from the extreme case in which $\varepsilon = 1.00$ and income effects are maximized. Keeping in mind that the two analyses differ in terms of various details (slight differences in each of time period considered, model specification and calibration procedure), we view this result as confirming that there is no inconsistency across the two studies. We conclude that the differing conclusions are due to the fact that the two papers document the effects of different changes in fundamentals. Indeed, Leonardi's and our exercises are complementary in the sense that the contribution obtained by the counterfactual proposed by Leonardi must be added to our numbers to obtain the total effect that is mediated via sectoral composition.⁴⁶

⁴⁶Because we do not explicitly target the moments highlighted by Leonardi (2015), one might interpret this comparison as a model validation exercise.

7 Sensitivity Exercises

In this section we report on four sensitivity exercises. In the first subsection we discuss the issue that our model does not contain investment and how controlling for this would affect our results. In the second subsection we discuss how allowing for trade would affect our findings. In the third section we discuss the implications of the possibility that observed changes in relative prices are biased upward due to mismeasured output. And in the fourth section we consider how our results are affected by allowing for a simple extension with endogenous skill supply.

7.1 Consumption Value Added vs Investment Value Added

Our analysis emphasizes the consequences of the systematic changes in the composition of output that accompany development for the overall demand for skill. Our model abstracts from investment and so implicitly focuses on systematic changes in the composition of consumption that accompany development. This raises the issue of whether the systematic changes in consumption value-added shares mimic the systematic changes in output value-added shares. Here we address this question and show that the two are very similar.

Herrendorf et al. (2013) carried out a similar exercise but for the traditional sectoral classification of agriculture, goods and services. They describe the procedure in detail in the online appendix to their paper. We follow their methodology but using our sectoral categories and so refer the reader to their paper for details.

Implementing the procedure in Herrendorf et al. (2013) uses the Historical I-O Tables produced by the BEA. Our analysis focuses on the period 1977-2005, and within this time period the I-O Tables are available at five year intervals from 1977 to 1997 and annually thereafter.

Before presenting the results we note two points. First, World KLEMS uses NAICS codes to define sectors, whereas the BEA uses SIC codes, so the sectoral comparisons are not perfectly matched. Second, the BEA and World KLEMS data sets use somewhat different measurement methodologies and as a result the levels of some variables will vary across the two. We have used the World KLEMS data for our output measures in order to have consistency with the labor measures that we use. But in view of these two issues, we are most interested in comparing the implications for the change in the value-added share of the skill-intensive sector.

When we carry out this exercise we find that the consumption value-added share for the skill-intensive service sector increases from 21% in 1977 to 36% in 2005. Our calibration exercise used output value-added shares from World KLEMS, and based on this data, we found that the output value added share of our service sector increased from 25% in 1977 to 39% in 2005.

While there are level differences between the two measurements, the key point for our purposes is that the increase in the share of the skill-intensive services sector is virtually identical between the two: 13.8% for the output value added share in our calibration exercise, versus 15.4% for the consumption value added share using the data from the BEA and the method of Herrendorf et al. (2013). We conclude that purging the data of investment is not an important concern.⁴⁷

This finding is perhaps not too surprising given the results in Herrendorf et al. (2020). They show that similar structural change has happened within both the consumption sector and the investment sector. In view of this it is not that surprising that the amount of structural change in consumption is similar to the amount of structural change found in total output.

7.2 Allowing for Trade

Our benchmark analysis considers a closed economy and so abstracted from changes in trade as a potential driving force. As we noted earlier, to the extent that much of trade takes place within the goods producing sector, it is possible that some of the skill-biased technological change that we infer reflects changes in composition within our low-skill intensive sector due to changes in specialization associated with trade. This alone would not affect our estimate of the contribution of the A_j to the overall change in the skill premium, though by diminishing the contribution of skill-biased technical change it would increase their relative importance.

More generally, lower trade costs can lead to greater specialization and hence higher productivity, so part of the productivity increases that we measure may result from trade. Our procedure aims to assess the contribution of productivity increases to the skill premium, but does not seek to understand the underlying source of the productivity increase. While we think it is of interest to assess the role of trade as

⁴⁷In fact, if we redo our calibration exercise using the consumption valued added shares from this calculation instead of the output value added shares that we originally used we find modestly larger effects. We now find that the contribution of the A_j is in the range of 19-26%.

a source of productivity growth, this issue is separate from the one we address. We refer the reader to the paper by Cravino and Sotelo (2019) for an analysis of the effects of lower trade costs on the skill premium in a framework similar to ours.

But not all trade takes place within the goods sector and the share of trade accounted for by trade in services is increasing over time. It is therefore possible that changes in trade patterns may also contribute to changes in the relative size of the skill-intensive sector. In this subsection we carry out a simple exercise to assess the potential magnitude of this effect. In particular, we will take sectoral net trade flows as given and solve for the equilibrium of our model given these flows.

Net sectoral trade flows create a wedge between production and consumption in each sector. If net exports from the skill-intensive sector are increasing over time, this would imply a decrease in consumption of the skill-intensive sector output holding labor allocations constant. Hence, this would create an incentive to increase the share of labor allocated to the skill-intensive sector in order to increase consumption from that sector. Similarly, if the imports of goods are increasing over time, then this would increase the relative consumption of low-skill intensive goods holding labor allocations fixed, and again create an incentive to reallocate labor to the skill-intensive sector. It follows that part of the movement of resources to the skill-intensive sector could be the result of changes in trade and not necessarily technology.

To estimate net trade flows for our two sector breakdown we do the following. From the Balance of Payments Accounts we obtain data on net trade flows for the “true” goods sector and the “true” services sector in the US economy for the full sample, 1977-2005. Over these years, the US ran a trade deficit in trade in “true” goods, and the deficit increased from around 1.4 percent of GDP to around 6 percent of GDP. Over the same time period the US ran a small trade surplus in “true” services, increasing from around 0.2 percent of GDP to around 0.5 percent of GDP. This trade surplus in services is to first approximation a trade surplus in skill-intensive services, as there is a small and relatively constant trade deficit in low-skill services, so that the overall change in the trade deficit in what we label the low-skill intensive sector (consisting of both goods and low-skill services) is to first approximation the same as the trade deficit coming purely from goods.⁴⁸ To evaluate these assumptions we can

⁴⁸Our net export figure for the goods sector is based on total value and is likely an overestimate of the deficit measured in terms of value added. For this reason we think our estimates for the effect of trade are likely an upper bound.

use disaggregated trade flows in services that are available from the BEA for a subset (1999-2005) of our simulation years. Splitting these flows into low- and high-skill intensive services components using our previous definitions and aggregate net trade flows to correspond to our model defined sectors, we show that our crude assumption, needed for the longer period, is a very good approximation over these years.

Taking net sectoral trade flows as given we implement the same calibration procedure as before and carry out the same counterfactuals to decompose the effects of technology. Intuitively, if net exports of the skill-intensive service sector are increasing over time, our calibration procedure would imply a lower value of \bar{c}_s , since the needed income effect from changes in technology would be reduced. Accordingly, the implied amount of skill-biased structural change would also be reduced.

The key message that results is that incorporating changes in trade has a relatively small effect on our results. In the interest of space we only report results for the benchmark case of $\varepsilon = .10$. Whereas our earlier results implied that changes in the A_{jt} 's accounted for between 18 and 24 percent of the overall change in the skill premium due to technical change, we now find that the range is between 16 and 21 percent. While the changing net sectoral trade balance does account for some of the movement of resources into the skill-intensive service sector, we find this effect to be relatively small.

7.3 Mismeasurement of Relative Price Changes

Here we consider the extent to which mismeasurement of relative prices might influence our results. Our quantitative analysis utilized information about changes in the relative price of the skill-intensive sector. Between 1977 and 2005 this relative price increased by more than forty-five percent. One possible concern is that price inflation in the skill-intensive sector might be upward biased because of the failure to properly account for quality improvements.

Here we report the results of a simple exercise to assess the extent to which our conclusions are affected by this possibility. In particular, consider the case in which the true increase in the relative price of the skill-intensive sector was only half as much as indicated by the official data. This means that real value added in this sector increased by roughly 30% more than indicated by the official data, and aggregate GDP grew by roughly 7 additional percentage points. Note that this adjustment has

no impact on the increase in the value added share of the skill-intensive sector.

We set $\rho = 1.53$ and $\varepsilon = 0.10$ and carry out the same calibration procedure as previously. Not surprisingly, given that we are holding ε fixed and decreasing the role of relative price changes, the calibration procedure yields a larger value for \bar{c}_S , indicating a larger role for nonhomotheticities. However, we find that the contribution of the skill-neutral component of technical change is virtually identical to what we found in our benchmark calculation. So while mismeasurement of relative price changes has implications for relative magnitudes of preference parameters, it has virtually no effect on our assessment of the role of demand factors. This follows naturally with the result that the different channels are less relevant to our quantitative impact as the overall amount of structural change.

7.4 Endogenizing the Supply of Skills

In the benchmark analysis we take the observed changes in the supply of skilled labor as an independent exogenous driving force. But if changes in the relative supply of skill are driven by changes in the skill premium, this specification may be inappropriate. Here we consider a simple extension in which changes in supply are completely due to changes in the skill premium to assess how this affects our results.

In particular, we assume a simple reduced form relationship between the supply of skill and the skill premium, $f_h = \bar{f}_h w_H^\zeta$. This simple relationship can be interpreted as a steady state supply function. We calibrate the parameters $\bar{f}_h = 0.15$ and $\zeta = 1.25$ so that we match the supply of skills in 1977 and 2005, when using the benchmark values for the preference and technology elasticities $\epsilon = 0.1$ and $\rho = 1.53$, respectively. We obtain a similar contribution for the changes in the A_{jt} 's in the range of 21–23%.

This result is perhaps not surprising. If changes in the supply of skill are completely driven by changes in the skill premium then changes in the supply of skill will have a similar decomposition and so will not affect the relative contribution of the different components of technology.

8 Cross-Country Analysis

In this section we extend our analysis to the ten other OECD countries for which the

available data exists and that we studied in Section 2.⁴⁹ The changes experienced by these countries differ quite significantly, both with regard to the change in the skill premium as well as the change in the relative supply of skilled labor. We view the results presented here as a simple first pass at extending the analysis to other countries.

In the interests of space we set $\varepsilon = 0.10$ for all countries and also set $\rho = 1.53$ as in our benchmark. We take as given the changes in the relative supply of skill in each country and infer country-specific processes for technical change using the same procedure described earlier. We choose country-specific values for a_G and \bar{c}_S to guarantee that the model generates the amount of structural change found in the data.⁵⁰

Using the country-specific calibrated models we carry out a decomposition exercise for each country corresponding to the results that we previously showed in Table 5. Results are in Table 8.

Table 8
Decomposition of Δw_H Due to Technology (%)

	ΔA_j Only	$\Delta \alpha_j$ Only	Interaction
Australia	3.9	86.0	10.1
Austria	21.1	71.8	7.1
Belgium	14.8	84.3	0.9
Denmark	10.0	85.7	4.3
Spain	20.7	75.6	3.7
Germany	21.8	76.9	1.3
Italy	21.3	53.1	25.6
Japan	11.4	83.3	5.3
Netherlands	19.6	79.4	1.0
UK	6.1	69.3	24.6
US	19.7	74.3	6.0
Median	19.6	76.9	3.5

The contribution of changes in the A_j alone varies significantly, from a low of 3.9%

⁴⁹Cross-country data is only available in EUKLEMS, so all of the results in this section use this data set. In particular, the results in this section for the US are based on EUKLEMS rather than World KLEMS, which explains why there are small differences from the earlier results. But as we emphasized earlier, the differences are quite minor.

⁵⁰In an earlier working paper version of this paper (Buera et al. (2015)) we showed that the implied series for technical change were quite similar across countries, which we think is reassuring.

in Australia to a high of 21.8% in Germany. But there is also considerable variation in the size of the interaction term across countries. If we allocate all of the interaction term to changes in the A_j then the range varies from a low of 14.0% (Australia and Denmark) to a high of 46.9% (Italy).

The key message from this brief examination of other countries is that the process of skill-biased structural change seems to play a significant role in many countries.

9 Conclusion

Using a broad panel of advanced economies, we have documented a systematic tendency for development to be associated with a shift in value added to skill-intensive sectors. It follows that development is associated with an increase in the relative demand for high-skill workers. We coined the term skill-biased structural change to describe this process. We have built a simple two-sector model of structural transformation and calibrated it to US data over the period 1977 to 2005 in order to assess the quantitative importance of this mechanism for understanding the large increase in the skill premium during this period. We find that technical change overall increased the skill premium by roughly 100 percentage points, and that between 18 and 24 percent of this change is due to the component of technical change that was sector-specific and skill-neutral, and that this component served to affect the skill premium through compositional changes. Moreover, this sector-specific skill-neutral component of technical change was also responsible for all of the structural change observed in the data.

Our findings have important implications for predicting the future evolution of the skill premium, since the continued growth of the value-added share of the skill-intensive sector will exert upward pressure on this premium even in the absence of skill-biased technological change.

In order to best articulate the mechanism of skill-biased structural change we have purposefully focused on a simple two-sector model. There is good reason to think that the mechanism we have highlighted is also at work at a more disaggregated level, so it is of interest to explore this mechanism in a richer model. It would also be of interest to explicitly introduce capital and investment into the analysis.

References

- ACEMOGLU, D. AND D. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” Elsevier, vol. 4 of *Handbook of Labor Economics*, chap. 12, 1043–1171.
- ACEMOGLU, D. AND V. GUERRIERI (2008): “Capital Deepening and Nonbalanced Economic Growth,” *Journal of Political Economy*, 116, 467–498.
- AGUIAR, M. AND M. BILS (2015): “Has Consumption Inequality Mirrored Income Inequality?” *American Economic Review*, 105, 2725–56.
- AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): “The Geography of Trade and Technology Shocks in the United States,” *American Economic Review*, 103, 220–25.
- BAUMOL, W. J. (1967): “Macroeconomics of Unbalanced Growth: Anatomy of an Urban Crisis,” *American Economic Review*, 57, 415–26.
- BERMAN, E., J. BOUND, AND Z. GRILICHES (1994): “Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures,” *The Quarterly Journal of Economics*, 109, 367–97.
- BERMAN, E., J. BOUND, AND S. MACHIN (1998): “Implications Of Skill-Biased Technological Change: International Evidence,” *The Quarterly Journal of Economics*, 113, 1245–1279.
- BOPPART, T. (2014): “Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences,” *Econometrica*, 82, 2167–2196.
- BOUND, J. AND G. JOHNSON (1992): “Changes in the Structure of Wages in the 1980’s: An Evaluation of Alternative Explanations,” *American Economic Review*, 82, 371–92.
- BUERA, F. J. AND J. P. KABOSKI (2009): “Can Traditional Theories of Structural Change Fit The Data?” *Journal of the European Economic Association*, 7, 469–477.
- (2012): “The Rise of the Service Economy,” *American Economic Review*, 102, 2540–69.

- BUERA, F. J., J. P. KABOSKI, AND R. ROGERSON (2015): “Skill Biased Structural Change,” *NBER Working Paper*.
- CASELLI, F. AND W. J. COLEMAN (2001): “The U.S. Structural Transformation and Regional Convergence: A Reinterpretation,” *Journal of Political Economy*, 109, 584–616.
- CERINA, F., A. MORO, AND M. RENDALL (2018): “The Role of Gender in Employment Polarization,” Manuscript.
- COMIN, D., D. LAKSHARI, AND M. MESTIERI (2015): “Structural Change with Long Run Income and Price Effects,” Manuscript.
- CRAVINO, J. AND S. SOTELO (2019): “Trade-Induced Structural Change and the Skill Premium,” *American Economic Journal: Macroeconomics*, 11, 289–326.
- DINARDO, J., N. M. FORTIN, AND T. LEMIEUX (1996): “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach,” *Econometrica*, 64, 1001–44.
- FEENSTRA, R. C. AND G. H. HANSON (1999): “The Impact Of Outsourcing And High-Technology Capital On Wages: Estimates For The United States, 1979-1990,” *The Quarterly Journal of Economics*, 114, 907–940.
- HALL, R. E. AND C. I. JONES (2007): “The Value of Life and the Rise in Health Spending,” *The Quarterly Journal of Economics*, 122, 39–72.
- HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2013): “Two Perspectives on Preferences and Structural Transformation,” *American Economic Review*, 103, 2752–89.
- (2014): “Growth and Structural Transformation,” in *Handbook of Economic Growth*, Elsevier, vol. 2 of *Handbook of Economic Growth*, chap. 6, 855–941.
- (2020): “Structural Change in Investment and Consumption? A Unified Analysis,” *The Review of Economic Studies*.
- JAEGER, D. E. (1997): “Reconciling the Old and New Census Bureau Education Questions: Recommendations for Researchers,” *Journal of Business & Economic Statistics*, 15, 300–309.

- KATZ, L. F. AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963-1987: Supply and Demand Factors,” *The Quarterly Journal of Economics*, 107, 35–78.
- KONGSAMUT, P., S. REBELO, AND D. XIE (2001): “Beyond Balanced Growth,” *Review of Economic Studies*, 68, 869–82.
- KRUSELL, P., L. E. OHANIAN, J. V. RIOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68, 1029–1054.
- LEONARDI, M. (2015): “The Effect of Product Demand on Inequality: Evidence from the US and the UK,” *American Economic Journal: Applied Economics*, 7, 221–247.
- MURPHY, K. M. AND F. WELCH (1992): “The Structure of Wages,” *The Quarterly Journal of Economics*, 107, 285–326.
- NGAI, L. R. AND B. PETRONGOLO (2014): “Gender gaps and the rise of the service economy,” LSE Research Online Documents on Economics 58152, London School of Economics and Political Science.
- NGAI, L. R. AND C. A. PISSARIDES (2007): “Structural Change in a Multisector Model of Growth,” *American Economic Review*, 97, 429–443.
- O’MAHONY, M. AND M. P. TIMMER (2009): “Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database,” *Economic Journal*, 119, 374–403.
- SWIECKI, T. (2017): “Determinants of Structural Change,” *Review of Economic Dynamics*, 24, 95–131.

Appendix 1

The results in Table 5 of the paper indicated that the contribution of changes in the A_j to the change in the skill premium were relatively constant across the four different specifications corresponding to different values for ε , the elasticity parameter in preferences. In this appendix we examine the reasons behind this result.

As noted in the text, the dominant channel through which changes in the A_j affect w_H is through changes in the relative sectoral expenditure shares. For this reason, we focus on why the contribution of the A_j to changes in sectoral expenditure shares is very similar across specifications.

Looking at this from the household perspective, the changes in expenditure shares are determined by the change in income and the change in relative prices. To understand the relative contributions of the three different driving forces (the f_i , the α_j and the A_j) to changes in relative expenditure shares it is therefore of interest to assess the effect of each driving force on income and relative prices.

It is useful to start with an accounting decomposition. Let π_I be the share of sectoral reallocation accounted for by the change in income (I) and let π_P be the share of sectoral allocation accounted for by changes in the price of services relative to goods (P), so that $\pi_I + \pi_P = 1$. Importantly, the value of the π_j will vary across the four specifications that have different values of ε . Next, let τ_{If} , $\tau_{I\alpha}$, and τ_{IA} denote the fraction of the overall change in income accounted for by the change in each of the f_i , the α_j and the A_j respectively. Similarly, let τ_{Pf} , $\tau_{P\alpha}$, and τ_{PA} denote the fraction of the overall change in relative prices accounted for by the change in each of the f_i , the α_j and the A_j respectively. By construction, $\tau_{If} + \tau_{I\alpha} + \tau_{IA} = 1$ and $\tau_{Pf} + \tau_{P\alpha} + \tau_{PA} = 1$. Thus, by construction, it follows that:

$$\pi_I(\tau_{If} + \tau_{I\alpha} + \tau_{IA}) + \pi_P(\tau_{Pf} + \tau_{P\alpha} + \tau_{PA}) = 1$$

The fraction of overall sectoral reallocation that is accounted for by changes in the A_j , which we denote by S_A , is then given by:

$$S_A = \pi_I \tau_{IA} + \pi_P \tau_{PA}$$

As noted above, the values of π_I and π_P vary across the two specifications, suggesting that S_A might also vary across specifications, though this will depend upon

the values of τ_{IA} and τ_{PA} and how they change across specifications.

Next we examine how each of the different subscripted π 's and τ 's vary across specifications. Table A1 shows the values for π_I and π_P .

Table A1
Values for π_I and π_P

	$\varepsilon = .01$	$\varepsilon = .10$	$\varepsilon = .50$	$\varepsilon = 1.00$
π_I	.25	.29	.50	1.00
π_P	.75	.71	.50	0.00

Table A2 reports the information necessary to compute the various subscripted τ 's. The first row of this table shows the total change in income and the relative price of services as we move from the initial equilibrium to the final equilibrium. By virtue of our calibration procedure, these changes are the same for each specification. The second row shows the effect on income and the relative price of services when we move from the initial equilibrium to the new equilibrium in which we change the value of the f_i from their initial value to their final value, holding the α_j and the A_j constant. The third and fourth rows show in turn the effect of changes in the α_j and the A_j , in each case starting from the equilibrium corresponding to the second row of the table in which the f_i take on their final values.

Table A2
Contributions to Changes in Income and Relative Prices

	$\varepsilon = .01$		$\varepsilon = .10$		$\varepsilon = .50$		$\varepsilon = 1.00$	
	ΔI	$\Delta \frac{p_s}{p_g}$	ΔI	$\Delta \frac{p_s}{p_g}$	ΔI	$\Delta \frac{p_s}{p_g}$	ΔI	$\Delta \frac{p_s}{p_g}$
Total Change	.184	.455	.184	.455	.184	.455	.181	.455
Change from Δf_i	.056	-.131	.056	-.131	.056	-.131	.54	-.132
Change from $\Delta \alpha_j$	-.060	-.004	-.060	-.005	-.059	-.011	-.058	-.022
Change from ΔA_j	.200	.631	.200	.631	.201	.635	.203	.646

The key takeaway from this table is that the contribution of each driving force to both margins (income and prices) is approximately constant across the four specifications, implying that the τ_{IA} and τ_{PA} are roughly constant. In particular, τ_{IA} is roughly constant at 1.09 and τ_{PA} is roughly constant at 1.38. As noted above, π_I varies from .25 to 1.00 as ε varies between .01 and 1.00. It follows that S_A varies

within the relatively small range of 1.30 to 1.09. The decrease from 1.30 to 1.09 represents a decrease of a bit more than 15%. This is consistent with the results displayed in Table 5, which show that the contribution of changes in the A_j to the change in w_H decreases from 18.1% to 15.3%, a drop of roughly 16% as we move from $\varepsilon = .01$ to $\varepsilon = 1.00$. Importantly, it is because the gap between τ_{IA} and τ_{PA} is not that large that the variation across specifications is relatively small.

The previous accounting decompositions raise two questions. First, why are τ_{IA} and τ_{PA} both roughly constant? And second, why is the gap between τ_{IA} and τ_{PA} not that large? We deal with each of these in turn.

We start with the first question, and proceed by considering each row in Table A2. The first row indicates that the effect of a change in the f_i on each of income and relative prices is roughly constant across the specifications. We argue that this is effectively a result of our calibration strategy. To see why, recall that in each specification we choose a value for ρ such that the implied effect of a change in the f_i on w_H is the same to first order. It follows that the effect of the change in the f_i on income are calibrated to be the same (to first order) in each specification. Additionally, the first order effect of a change in the f_i on relative prices is only a function of the factor shares. Because our calibration procedure necessarily matches the factor payment shares in the data, it follows that changes in relative prices are also the same to first order. In summary, the fact that the effect of changes in the f_i holding the α_j and the A_j constant on each of income and relative prices is roughly constant across the four specifications is a feature of our calibration strategy.

Next we consider the effects due to a change in the α_j . First we consider the effects on income. To do this we ask how a change in α_j affects output in sector j . To first order, one can show that the effect is given by:

$$\frac{dy_j}{y_j} = \frac{\rho}{\rho - 1} \frac{(\theta_j - \alpha_j)}{(1 - \alpha_j)} \frac{d\alpha_j}{\alpha_j}$$

Holding ρ fixed, the effect of a given percentage change in α_j on y_j is dictated by factor shares. Our calibration procedure necessarily matches the factor shares in all specifications, so that this effect will be the same across specifications. However, as noted above, the value of ρ is not constant across specifications, and because the value of ρ affects the imputation of the changes in the α_j , it follows that both the value of ρ and the percent change in the α_j will vary across the specifications as ε

varies. But we argue that these effects are small. First, as ε varies from .01 to .50, the implied changes in ρ are second order. But second, and more generally, our procedure for computing the series for the α_j imposes that the following first order condition holds at each point in time:

$$\frac{H_j}{L_j} = \left(\frac{1}{w_H}\right)^\rho \left(\frac{\alpha_j}{1 - \alpha_j}\right)^\rho$$

That is, any change in the value of ρ will lead to a new series for $\alpha_j(1 - \alpha_j)$ that keeps the value of $[\alpha_j/(1 - \alpha_j)]^\rho$ constant. And because this value remains constant across specifications it follows that the effect on relative labor demands will also be constant across specifications. Put somewhat differently, any change in the value of ρ will be offset by a corresponding change in the implied series for $\alpha_j(1 - \alpha_j)$ that is neutral with respect to the impact on relative demands for labor. We conclude that the effect of the change in the α_j on income is roughly constant largely on account of our calibration procedure.

To first order, the effect of the changes in the α_j on the relative price of services depends on how the changes in the α_j affect relative costs of production across the two sectors. A given increase in each of the α_j will have a larger impact on the costs in the service sector given that high skill labor is a higher fraction of total costs in that sector. However, our procedure implies that the increase in α_G is larger than the increase in α_S , and roughly offsets the first factor just mentioned. It follows that the effect of the change in the α_j on relative prices is effectively zero. Across specifications, we change both the implied changes in the α_j as well as the effect of a given change in the α_j on costs. These effects are roughly offsetting, which is why the overall result is effectively constant across specifications. In summary, the fact that this effect is zero is a feature of the data. Had the changes in the two α_j been different, the overall effect would have not been zero.

To this point we have argued that the constancy of the effects of changes in the f_i and the α_j are largely a feature of our calibration strategy and the procedure for measuring the α_j . If these effects are independent of the value of ε and the total effect on income and prices is the same for all values of ε , it follows that the effect of the change in the A_j must also be constant. So the answer to the first question, concerning the relative constancy of τ_{IA} and τ_{PA} is that this is effectively an implication of our calibration strategy, in that for each value of ε , the values of

ρ and \bar{c}_s are adjusted so that both the pure effect of a change in the f_i on w_H is constant across specifications, that the total effect on both income and relative prices is the same across all specifications, and that factor and expenditure shares are all the same across all specifications.

Given that the values of τ_{IA} and τ_{PA} are roughly constant and that π_I varies from .25 to 1.00, the variation in S_A will be dictated by the gap between τ_{IA} and τ_{PA} . As noted above, we find that this gap is roughly .30, so that the variation in S_A is roughly .20 as π_I varies from .25 to 1.00. The fact that the gap is equal to .30 is a feature of the data and not a feature of our calibration.