

Skill-Biased Structural Change*

Francisco J. Buera
Federal Reserve Bank of Chicago

Joseph P. Kaboski
University of Notre Dame

Richard Rogerson
Princeton University

May 3, 2015

Abstract

We document for a broad panel of advanced economies that increases in GDP per capita are associated with a shift in the composition of value added to sectors that are intensive in high-skill labor. It follows that further development in these economies leads to an increase in the relative demand for skilled labor. We develop a two-sector model of this process and use it to assess the contribution of this process of skill-biased structural change to the rise of the skill premium in the US, and a broad panel of advanced economies, over the period 1977 to 2005. We find that these compositional demands account for between 25 and 30% of the overall increase of the skill premium due to technical change.

*We thank Daron Acemoglu, David Dorn, Chad Jones, Pete Klenow as well as seminar participants at the 2015 AEA Meetings, Pittsburgh, Stanford, and USC for useful comments.

1 Introduction

The dramatic 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 in recent decades. Understanding the underlying causes of these trends, both the common features and those unique to different country’s experiences, is important for projecting future trends and evaluating the extent to which policies might be effective or advisable in addressing this issue. It has therefore been a major research focus. The consensus view in the literature, as summarized in the review article by Acemoglu and Autor (2011), is that the dominant factor generating this trend is skill-biased technological change (SBTC).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 high-skill intensive industries that accompanies the process of continued development among advanced economies.

The economic rationale for the consensus view is based on simple demand and supply analysis in the context of what Acemoglu and Autor call the canonical model—an aggregate production function that has high-skilled and low-skilled labor as inputs. Because the relative supply of high-skilled labor has increased, an increase in the relative wage of high-skilled labor requires some force that increases the relative demand for high-skilled labor. But the only way to generate an increase in relative demand for high-skilled labor in the canonical model is via technical change that favors high-skilled workers.²

The economic intuition behind our result is equally simple, but it requires that one go beyond the aggregate production function approach in the canonical model. 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 this process alone will increase the demand for high-skilled workers, independently of whether it is driven by skill-neutral or skill-biased technical change. With an aggregate production function, development that comes from skill-neutral technical change has no effect on the relative demand for high-skilled workers, so this channel is absent.

To assess the quantitative significance of this channel, we develop a simple general equilibrium model of structural transformation that incorporates an

¹This 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.

²One can formulate richer structures in which the fundamental technical change is not explicitly skill biased but which operates like skill-biased technical change in a reduced form sense. For example, Krusell et al. (2000) stressed how the decreasing price of equipment coupled with capital-skill complementarity also serves to increase the relative demand for skill.

important role for skill and use it to study the evolution of the US economy between 1977 and 2005. In order to best highlight the shift in value added to high skill-intensive sectors, we study a two-sector model in which the two sectors are distinguished by their intensity of skill workers in production. We allow for sector-specific technological change, which is a (sector-specific) combination of skill-neutral and skill-biased technical change. We show how the model can be used to infer preference parameters and the process for technical change using data on the change in the composition of employment by skill, the change in aggregate output, changes in sectoral factor shares, the skill premium, relative sectoral prices and the distribution of sectoral value added.

In the data, our measure of the skill premium increases from 1.41 to 1.90 between 1977 and 2005, an increase of 49 percentage points.³ Our calibrated model perfectly matches this increase. We then use the model to decompose this increase into three different components: one due to the changes in the relative supply of high-skill workers, one part that is due to skill-biased technical change, and a third part due to other technological changes. If there had been no change in technology, our model predicts that the skill premium would have decreased to 0.92, a drop of 49 percentage points, due to the increase in the relative supply of high-skill workers. It follows that overall changes in technology created an increase in the skill premium of 98 percentage points. In our benchmark specification, between one quarter and one third of this increase comes from changes in technology other than skill-biased technical change, operating through their effect on the composition of value added. We conclude that systematic changes in the composition of value added associated with the process of development are an important factor in accounting for the rise in the skill premium. In fact, if skill-biased technical change had been the sole source of technical change over this period, our model predicts that the skill premium would have increased by only 18 percentage points instead of by 49.

Having established the importance of this effect for the US, we also repeat the analysis for a set of nine other OECD countries. While there is some variation in the contribution of compositional changes in value added to changes in the skill premium across countries, ranging from around 10 percent to slightly more than 50 percent, the average for this sample is slightly above 30 percent, very much in line with our estimates for the US.

Our paper is related to many papers in two distinct literatures, one on SBTC and the skill premium and the other on structural transformation. An early contribution in the former literature is Katz and Murphy (1992), and an excellent recent overview is provided by Acemoglu and Autor (2011). Although Katz and Murphy predominantly use an aggregate production function to interpret the data, Section V of their paper does argue that changes in sectoral composition are an important element of the increased relative demand for skill. Except for their explicit consideration of international trade, they are otherwise agnostic about the driving forces behind these reallocations. Relative to them, our

³Our measure of the skill premium compares those with *at least* college degrees to those with high school degrees *or less*, and is based on total compensation and not just wages and salaries, which explains why this increase is larger than what the literature typically reports.

contribution is threefold. First, we document the importance of compositional effects that are systematically related to the process of development. Second, we show how to uncover the different dimensions of technological change in a multi-sector framework. Third, we present a general equilibrium model in which one can assess the driving forces behind compositional changes. It is also of interest to note that whereas their analysis ended in 1987, ours covers the period through 2005. 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.

The paper that we are most closely related to is Buera and Kaboski (2012). Like us, they study the interaction between development and the demand for skill, though their primary contribution is conceptual, building a somewhat abstract model to illustrate the mechanism. Relative to them our main contribution is to build a simple model that can easily be connected to the data and to use the model to quantitatively assess the mechanism. Leonardi (forthcoming) considers a similar mechanism to us, but focuses on how demand varies by education attainment as opposed to income more broadly, and finds relatively small demand effects.⁴ An important antecedent of our work is the paper by 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 differently than us, they focus on differential intensities for physical capital and the role of the relative price of physical capital rather than human capital. Their work is also primarily theoretical.

An outline of the paper follows. Section 2 presents aggregate evidence on the relation between development and the value added share for high skill intensive services in a panel of advanced economies, in addition to some other important empirical patterns. Section 3 presents our general equilibrium model 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 contribution of various factors to the evolution of the skill premium. Section 6 assesses the contribution of skill-biased structural change for relative prices, and in Section 7 we extend our analysis to a set of nine other countries. Section 8 concludes.

⁴In independent work, Boppart (2015) carries out a similar exercise to ours for the US. Ngai 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.

2 Empirical Motivation

This section provides motivating facts for the prevalence of what we refer to as skill-biased structural change. In particular, using data for a broad panel of advanced economies, we document two key facts. First, there is a strong positive correlation between the level of development in an economy, as measured by GDP per capita, and the share of value added that is attributed to high skill services. Second, there is also a strong positive correlation between the level of development and the price of high skill services relative to other goods and services. Interestingly, these relationships are very stable across countries, and in particular, the experience of the US is very similar to the average pattern found in the data.

We supplement the above aggregate time series evidence for a panel of countries with some evidence about cross-sectional expenditure shares in the US economy. In particular, we show that the expenditure of higher income households contains a higher share of skill-intensive value-added. This fact will serve two purposes. First at a qualitative level it is suggestive evidence of the role of a non-homotheticity in the demand for high skill services, which is a feature we will include in our model. Second, we will also show how this cross-sectional moment provides important information about preference parameters that is not readily available from aggregate time series data.

2.1 Aggregate Panel Evidence

The starting point for our analysis is the earlier work of Buera and Kaboski (2012). They divide industries in 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 group rose substantially between 1950 and 2000, the value added share of the low-skill group actually fell over the same time period. This finding suggests that the traditional breakdown of economic activity in the structural transformation literature, into agriculture, manufacturing and services, was perhaps not well suited to studying the reallocation of economic activity in today’s advanced economies. Here we pursue this line of work further, modifying their aggregation procedure to include goods-producing industries, and extending their analysis to a broad panel of advanced economies.

The analysis is based on underlying value added data from the EUKLEMS Database (“Basic Table”)⁵ These data exist in comparable form for a panel of 12 advanced economies over the years 1970-2005.⁶ The sectoral value-added data are available at roughly the 1 to 2-digit industry level. We focus on a two-way split of industries into high skill intensive and low skill intensive based on the

⁵See O’Mahony and Timmer (2009).

⁶These countries are Australia, Austria, Denmark, France, Germany, Italy, Japan, the Netherlands, South Korea, Spain, the United Kingdom, and the United States. The U.S. data for value-added go back to only 1977, while the Japan data go back to only 1973.

share of labor income paid to high-skill workers.⁷ While one could imagine more detailed splits, including more than two skill categories and perhaps interacting skill intensity with goods vs. services, we feel that this two-way split both facilitates exposition and allows us to capture a robust pattern in the cross-country data.

The labor payment data come from the EUKLEMS Labour Input Data and are slightly more disaggregated. High skill-intensive service sectors are: “Financial Intermediation”, “Real Estate and Business Services”, “Education”, and “Health and Social Work”. In 1970, the economy-wide average share of labor compensation paid to high-skill workers in the U.S. was 20 percent; the corresponding shares for these high skill-intensive industries were 34, 38, 74, and 49 percent, respectively. These industries remain well above average throughout the time period.⁸ We combine these data with real (chain-weighted) GDP per capita data from the Penn World Tables 7.1. Finally, we demean both the value-added share data and the (log) GDP per capita data by taking out country fixed effects.




Figure 1: Value-added by Skill-intensity and Economic Development.

Figure 1 shows the data pooled across time and countries. The small squares show the relationship for the panel of advanced countries; we have highlighted the United States data using the larger circles. The relationship is clear: the

⁷High-skill is defined as a college graduate and above.

⁸The next highest industries are “Chemicals and Chemical Products” (27 percent), “Coke, Refined Petroleum, and Nuclear Fuel” (21 percent), and “Electrical and Optical Equipment” (21 percent).

value added share of the high skill-intensive sector increases with log GDP/capita, with a highly significant (at a 0.1 percent level) semi-elasticity of 0.17. The regression line implies an increase of roughly 24 percentage points as we move from a GDP per capita of 10,000 to 40,000 (in 2005 PPP terms), and explains 80 percent of the variation in the data. Moreover, we see that the United States data is quite similar to the overall relationship. Indeed, the tight relationship suggests that 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 sum, the tendency for economic activity to move toward high skill-intensive services as an economy develops is a robust pattern in the cross-country data.

One of the common explanations for structural change is changes in relative prices. (See, for example, Baumol (1967) and Ngai and Pissarides (2007).) Using value-added price indices from the same EUKLEMS Database, we can examine the correlation between the relative price of high-skill intensive services and the increasing value added share of high-skill intensive services that accompanies the process of development.⁹ Figure 2 is analogous to Figure 1, but it plots the price index of the high skill-intensive sector relative to the low-skill intensive sector rather than share data on the y-axis. Again we have demeaned both the relative price and log GDP per capita data to eliminate country fixed effects, and normalized the relative price indices to 100 in 1995. As before, the larger circles represent the U.S. data.




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

⁹We construct sector-level aggregate indices as chain-weighted Fisher price indices of the price indices for individual indices. Calculation details are available in the data appendix.

Again, the relationship is striking. The linear regression is highly significant, explains 84 percent of the variation in the demeaned data, and is quantitatively important: the relative price of the high skill-intensive sector increases almost two and a half times over the range of the data. Finally, the U.S. relationship is quite similar to the overall relationship, and again the tight relationship suggests that cross-country variation in this relative price-income relationship is second order. We conclude that changes in relative prices are another robust feature of the structural transformation process involving the movement of activity toward the high-skill intensive sector.

2.2 Income Effects: Cross-Sectional Household Evidence

A second common explanation for structural change is income effects associated with non-homothetic preferences. (See, for example, Kongsamut et al. (2001).) With this in mind it is of interest to ask whether high-skill intensive services are a luxury good, i.e., have an income elasticity that exceeds one. To pursue this we examine whether the relationship between the skill intensity of value-added consumption and income exists in the expenditure data from a cross-section of households. 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.

One complication with pursuing this approach is that it involves mapping household expenditure data through the input-output system in order to determine the consumption shares of value added. We briefly sketch the steps of this procedure here, and provide more details in the appendix. First, we start with the household level Consumer Expenditure Survey (CEX) data for the United States from 2012. We adapt a Bureau of Labor Statistics (BLS) mapping from disaggregate CEX categories to 76 NIPA Personal Consumption Expenditure (PCE) categories. We 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. Using the EUKLEMS data, we classify the 69 industries as high skill- or low skill-intensive.¹⁰ We therefore have constructed a household-level data set of the amount of value added per dollar spent produced by high skill-intensive sectors and low skill-intensive sectors, which we can regress on household observables, most importantly income or education, and potentially a host of other household level controls. We restrict ourselves to the primary interview sample, and each observation is a household-month observation.

¹⁰The labour data from EUKLEMS contains 41 distinct industries. The "basic" data, from which we obtain value-added data, contain only 33 distinct industries.

Table 1 presents results for regressions of the total share of expenditures that is high-skill intensive. 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 share of value-added embodied in expenditures is 0.012.

Table 1: **Household High-Skill Intensive Expenditure Share vs. Income**

| | OLS | IV | OLS |
|------------------------|----------|----------|----------|
| Ln Income | 0.012*** | 0.049*** | . |
| SE | 0.001 | 0.002 | . |
| High Skill Head | . | . | 0.043*** |
| SE | . | . | 0.002 |
| R^2 | 0.08 | 0.02 | 0.15 |
| Observations | 48,550 | 48,550 | 17,812 |

*** indicate significance at the 1 percent level.

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 skilled is defined as 16 years of schooling attained, while low skilled is defined as 12 years attained.

Income is certainly mis-measured in the micro data, and even if properly measured it would only proxy for permanent income, leading to a likely attenuation bias. The second column attempts to alleviate this measurement error by instrumenting for log income using the years of schooling attained by the head of household. Instrumenting for income in this fashion increases the coefficient roughly four-fold to 0.049.

The last column uses education as a direct regressor, replacing log income with a dummy for whether the head of household is high skilled 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 from the sample.) The coefficient indicates that the share of value-added embodied in expenditures is 4.3 percentage points higher in households with a high-skilled head.

We have examined the robustness of the results in Table 1 to alternatives. Table 1 uses monthly expenditures of the household, but if we aggregate household expenditures across the three months they are surveyed, we find nearly identical results. 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, but the raw estimates are similar (0.032 rather than 0.043).¹¹

¹¹An alternative analysis, however, where we try to directly estimate expenditure elastic-

Second, dropping demographic controls increases the sample by about 15 percent and lowers the coefficients on income somewhat (by roughly 50 percent), but the coefficients remain highly significant. Dropping the controls have essentially no impact on the high-skilled head of household coefficient. The main effect of dropping the controls is substantially lower R^2 values.

Quantitatively, even the larger, instrumented, income coefficient of 0.049 is substantially smaller than the aggregate time series value of 0.17 in Figure 1, but not negligible in comparison. We therefore take this as suggestive 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.

While we do regard this evidence as suggestive, we want to 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 insurance premiums (a substantial share of which is paid by employers) and higher education (a substantial share of which is paid by government).¹²

2.3 Summary

In summary, we have documented a robust relationship in the time series data for advanced economies regarding the movement of activity into high-skill intensive services and the process of development. We will refer to this process as skill-biased structural transformation, so as to emphasize both its connection to the traditional characterization of structural transformation 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 transformation literature emphasizes the role of both non-homotheticities and relative price changes as drivers of structural transformation, and we have also presented evidence that both of these effects are relevant in the context of skill-biased structural transformation as well. Specifically, we documented a strong positive correlation between the relative price of high skill intensive services and GDP per capita in a cross-country panel as well as a positive correlation between household value added expenditure shares on high skill intensive services in the US cross-section. These relation-

ities by regressing the (log) level of high-skill value-added on the (log) level of expenditures gives substantially higher estimates. This is driven by certain lumpy expenditures like higher educational expenses and car purchases driving both up in particular months. This motivates an emphasis on the relationship between the high skill share and education or income rather than direct expenditures per se.

¹²The 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.

ships are not only highly statistically significant, but they are also economically significant in a quantitative sense.

Finally, we should note that in documenting these relationships we have used a strict high versus low-skill dichotomy. This masks important within sector heterogeneity. Indeed, within the low-skill sector, a pattern emerges that the relatively more skill-intensive sectors, e.g., manufacturing industries like electrical equipment and chemicals expand relative to the less skill-intensive sectors like agriculture or textiles.¹³ In this sense, our simple dichotomy may understate the true extent of skill-biased structural change. However, the relative price patterns, use patterns (consumption and investment), and trade patterns make the analysis at a more disaggregated level more difficult to interpret and much less directly tied to traditional structural change forces.

3 Model and Equilibrium

Our analysis focuses on how features of intratemporal equilibrium allocations are affected by changes in the economic environment over time that operate through changes in income and relative prices. In order to capture these interactions in a simple a setting as possible, our benchmark model is static. In this section, we describe the economy and its equilibrium at a point in time; later we describe the features that we will allow to change over time to generate skill-biased structural change as described in the previous section. 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.

3.1 Model

There are two types of households in the economy, distinguished by their skill level. We will refer to them as low skilled and high skilled, denoted by the subscript L and H respectively. The total mass of households is normalized to one and we denote the fractions of low and high skill households as f_L and f_H respectively, where $f_L + f_H = 1$. There are multiple sectors in the economy, distinguished by the extent to which they rely on low versus high-skilled labor. To facilitate the analysis we assume that there are only two sectors, which in the calibration of the model will be connected to the low and high-skilled aggregates studied in the previous section. As a practical matter, most high-skilled sectors are services and most goods sectors are low-skilled sectors. For this reason it will be convenient to label the two sectors as goods and services even though these labels are not strictly correct. This will allow us to avoid having double indices to distinguish between both low- vs. high-skilled workers and low vs. high skill-intensive sectors.

¹³Katz and Murphy (1992) give a detailed analysis across 150 2-digit industry-occupation cells for the period, 1963-1987. Autor and Dorn (2013) present a recent account focusing on detailed occupation categories.

All households have identical preferences over goods and services, independently of their skill level. We assume these preferences take the form:

$$a_G c_{Gi}^{\frac{\varepsilon-1}{\varepsilon}} + (1 - a_G) (c_{Si} + \bar{c}_S)^{\frac{\varepsilon-1}{\varepsilon}}$$

where c_{Gi} and c_{Si} are consumption of goods and services by an individual of skill level i , $0 < a_G < 1$, $\bar{c}_S \geq 0$ 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.¹⁴ This non-homotheticity will be important in allowing the model to match the observation from the previous section concerning development and the value-added share for high-skilled sectors. Note that households are assumed to not value leisure, since our focus will be on the relative prices of labor given observed supplies.

Each of the two production sectors has a constant returns to scale production function that uses low- and high-skilled labor as inputs. We assume that each of these production functions is CES:

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

where L_j and H_j are inputs of low- and high-skilled labor in sector j , respectively. The parameter α_j will dictate the importance of low- versus high-skilled labor in each sector. While one could imagine that the elasticity of substitution between these two factors also differs across sectors, our benchmark specification will assume that this value is the same for both sectors. We consider the effects of cross-sectional variation in this parameter in our sensitivity analysis.

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-skilled 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-skilled labor to unity so that the price of high-skilled labor will also represent the skill premium, though in our derivations it will be convenient to postpone implementing this normalization.

The definition of competitive equilibrium for this model is completely standard and straightforward, so here we will focus on characterizing the equilibrium.

¹⁴This 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 (2014) 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.

Individuals of skill $i = L, H$ solve

$$\max_{c_{Gi}, c_{Si}} a_G c_{Gi}^{\frac{\varepsilon-1}{\varepsilon}} + (1 - a_G) (c_{Si} + \bar{c}_S)^{\frac{\varepsilon-1}{\varepsilon}}$$

subject to

$$p_G c_{Gi} + p_S c_{Si} = w_i. \quad (1)$$

The first-order conditions associated with their problem are

$$a_G c_{Gi}^{-\frac{1}{\varepsilon}} = p_G \lambda_i \quad (2)$$

$$(1 - a_G) (c_{Si} + \bar{c}_S)^{-\frac{1}{\varepsilon}} = p_S \lambda_i \quad (3)$$

Rearranging these two first order conditions and substituting into (1) yields:

$$c_{Gi} = \frac{w_i + p_S \bar{c}_S}{p_S \left(\frac{1-a_G}{a_G} \frac{p_G}{p_S} \right)^\varepsilon + p_G}$$

and

$$c_{Si} = \frac{\left(\frac{1-a_G}{a_G} \frac{p_G}{p_S} \right)^\varepsilon w_i - p_G \bar{c}_S}{p_S \left(\frac{1-a_G}{a_G} \frac{p_G}{p_S} \right)^\varepsilon + p_G}$$

Normalizing w_L to unity the aggregate expenditure share for services is then given by:

$$\frac{p_S [(1 - f_H) c_{SL} + f_H c_{SH}]}{1 - f_H + f_H w_H} = \frac{p_S \left(\frac{1-a_G}{a_G} \frac{p_G}{p_S} \right)^\varepsilon}{p_S \left(\frac{1-a_G}{a_G} \frac{p_G}{p_S} \right)^\varepsilon + p_G} - \frac{p_S}{1 - f_H + f_H w_H} \frac{p_G}{p_S \left(\frac{1-a_G}{a_G} \frac{p_G}{p_S} \right)^\varepsilon + p_G} \bar{c}_S \quad (4)$$

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

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

The first order conditions are

$$p_j A_j \left[\alpha_j H_j^{\frac{\rho-1}{\rho}} + (1 - \alpha_j) L_j^{\frac{\rho-1}{\rho}} \right]^{\frac{1}{\rho-1}} (1 - \alpha_j) L_j^{-\frac{1}{\rho}} = 1 \quad (5)$$

$$p_j A_j \left[\alpha_j H_j^{\frac{\rho-1}{\rho}} + (1 - \alpha_j) L_j^{\frac{\rho-1}{\rho}} \right]^{\frac{1}{\rho-1}} \alpha_j H_j^{-\frac{1}{\rho}} = w_H \quad (6)$$

Taking the ratio of (5) and (6) we obtain

$$\frac{H_j}{L_j} = \left(\frac{\alpha_j}{1 - \alpha_j} \frac{1}{w_H} \right)^\rho \quad (7)$$

Substituting (7) into (5) and rearranging delivers an equation for the price of sector j output in terms of the skill premium w_H :

$$\hat{p}_j(w_H) = \frac{1}{A_j} \left[\frac{\alpha_j^\rho}{w_H^{\rho-1}} + (1 - \alpha_j)^\rho \right]^{\frac{1}{1-\rho}}. \quad (8)$$

The above expression implies that the search for equilibrium prices can be reduced to a single dimension: if we know the equilibrium wage rate for high-skilled labor then all of the remaining prices can be determined.

In what follows we derive an expression for the market-clearing condition for high-skilled labor that contains the single price w_H . Using (7), the production function of sector j , and (8), we obtain

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

Similarly, we solve for the demands for high-skilled labor in sector j as a function of output of sector j and the wage for high-skilled labor:

$$H_j = \left[\frac{\alpha_j \hat{p}_j(w_H) A_j}{w_H} \right]^\rho \frac{Y_j}{A_j} \quad (9)$$

Using (9) and equilibrium in the goods market, the market-clearing condition for high-skilled labor yields a single equation in w_H :

$$\begin{aligned} & \left[\frac{\alpha_S \hat{p}_S(w_H) A_S}{w_H} \right]^\rho \frac{\sum_{i=L,H} f_i \hat{c}_{Si}(w_H)}{A_S} \\ & + \left[\frac{\alpha_G \hat{p}_G(w_H) A_G}{w_H} \right]^\rho \frac{\sum_{i=L,H} f_i \hat{c}_{Gi}(w_H)}{A_G} = f_H. \end{aligned} \quad (10)$$

where 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 $\hat{p}_i(w_H)$ defined in (8).

4 Accounting for Growth and Structural Transformation

In this section we describe how the model of the previous section can be calibrated so as to account for growth, structural transformation, and the skill premium. To do this we will use the above model to explain equilibrium outcomes at two different points in time, that we denote as 0 and T for the initial and terminal periods respectively. We assume that there are two types of changes in the economic environment across these two periods. One is a change in the fraction of individuals that are high skill. The other is technological change. However, we will allow for technological change along several dimensions. In

particular, we assume that technological change within each sector can be some combination of skill biased technological change and skill neutral technological change. Skill-biased technological change in sector j will be captured by changes in α_j , whereas neutral technological change will be captured by changes in A_j . Additionally, the mix of these two types of technological change is allowed to differ across the two sectors. 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 do not allow for preferences to change over time.

In choosing our targets, we again focus on the KLEMS data from Section 2. For the U.S., the complete data are available for the years, 1977 to 2005, so those two end years become our targets for 0 and T .¹⁵ These dates are convenient, since 1977 effectively marks a local minimum in the skill premium (see Acemoglu and Autor (2011) for earlier data), and it secularly increases after 1977. The data contain total compensation and total hours by industry, skill level (“low”, “medium”, and “high”, which are effectively, secondary degree or less, some tertiary schooling, and four year college degree or more), gender, and age groupings (15-29, 30-49, and 50 and over). We combine the compensation of KLEMS categories of “low” and “medium” educated workers of all genders and ages into our classification of low-skilled, and all varieties of “high” into high-skilled, in order to calculate aggregate and goods and service sector-specific labor income shares, using the same sectoral classification as in Section 2. To compute the skill premium, we compute wages as the ratio of compensation to hours, and the skill premium as the ratio of college-educated (“high”) to high school-educated (“low”) prime-aged (i.e., aged 30-49) male wages. This premium rises from 1.41 in 1977 to 1.90 in 2005.¹⁶ Finally, we infer f_H , using the identity that the ratio of labor compensation equals the product of the skill premium and the relative quantity of high- to low-skilled labor (f_H and $f_L = 1 - f_H$, respectively). In doing so, the implicit assumption is that relative wages between different low-skilled (high-skilled) demographic groups reflect relative efficiency units of low-skilled (high-skilled) labor.

In mapping the model to these targets, we proceed in two steps. In the first step, we describe a procedure that allows us to infer the changes in parameters that represent technical change given a value for ρ . In the second step, we describe how to infer values for the preference parameters.

We begin with the determination of technological change. First, we show

¹⁵The KLEMS value-added and price data line up quite closely with BEA data, but the consistent aggregation is not available prior to 1977. They are available through 2007, but labor compensation and hours data are only available through 2005.

¹⁶Comparing earnings of full time workers using CPS data, Figure 1 in Acemoglu and Autor (2011) indicates values of 1.39 and 1.64 for 1977 and 2005 respectively. Our measure indicates a roughly 25 percentage point greater increase. This difference reflects two factors. First, our lower skill group includes those with less than high school education in addition to those with high school education, while our high skill group involves those with *at least* a college degree. Second, our measure is based on total compensation and not just on labor earnings. Pierce (2010) documents that the change in the 90-10 ratio over this time period is more than twenty log points higher when using total compensation than when using CPS wages. His analysis is based on firm level data and so does not allow a breakdown by educational attainment.

that 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, note that the share of sector j income going to high skill labor is

$$\begin{aligned}\theta_{Hjt} &= \frac{w_{Ht}H_{jt}}{\hat{p}_j(w_{Ht})Y_{jt}} \\ &= \frac{\alpha_{jt}^\rho}{\alpha_{jt}^\rho + (1 - \alpha_{jt})^\rho w_{Ht}^{\rho-1}}\end{aligned}$$

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

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

Next we turn to determining the values of the A_{jt} 's. Although there are four of these parameters, the two values in period 0 basically reflect a choice of units and so can be normalized, leaving only two parameters to be determined. We next derive a condition that determines the two A_{jT} 's up to a scale factor. As is well known in the literature, if we had Cobb-Douglas sectoral technologies with identical labor share parameters, then relative sectoral prices would simply be the inverse of relative sectoral TFPs. In this case, changes in relative prices would pin down changes in relative TFPs, and hence determine the values of the two A_{jT} 's up to a scale factor. This same relation holds more generally, and in particular would also apply if the sectoral production functions are CES with identical parameters. While this result does not apply to our setting because of sectoral heterogeneity in the α_{jt} 's, there is still a close connection between relative sectoral prices and relative sectoral TFPs (i.e., the A_{jt}), though the relation is somewhat more complex. In particular, using equation (8) for the two sectors we have:

$$\frac{A_{Gt}}{A_{St}} = \frac{p_{St}}{p_{Gt}} \left[\frac{\frac{\alpha_{Gt}^\rho}{w_{Ht}^{\rho-1}} + (1 - \alpha_{Gt})^\rho}{\frac{\alpha_{St}^\rho}{w_{Ht}^{\rho-1}} + (1 - \alpha_{St})^\rho} \right]^{1/(1-\rho)}. \quad (11)$$

As noted above, the initial values of the A_{jt} simply reflect a choice of units. We will normalize A_{S0} to equal one, and given the calibrated values for the α_{j0} and the value of w_{H0} , we will choose A_{G0} so as to imply $p_{G0}/p_{S0} = 1$. In this case p_{GT}/p_{ST} can be easily identified with the change in the relative sectoral prices, and this value from the data will pin down the sectoral TFPs in period T up to a scale factor.

This scale factor will of course influence 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. However, in order to compute aggregate output at a point in time (and thus also the growth rate in aggregate output) it is necessary to determine the sectoral distribution of output. The

relations that we have imposed thus far guarantee that maximum profits will be zero in each sector but do not determine the scale of operation in either sector. Intuitively, the split of activity across sectors at given prices will be determined by the relative demands of households for the two sectoral outputs at those same prices. Below we will describe how preference parameters can be chosen to match the sectoral distribution of value added at both the initial and final date. For now we simply note that if we assume this split is the same as in the data, then given our previous steps we can compute the growth rate of output per worker as a function of the scale factor and hence use the growth rate in aggregate output to determine the scale factor.

To summarize, at this point, given a value for ρ , we have identified all of the technology parameters given values for f_{Lt} at $t = 0, T$, θ_{jt} for $j = G, S$, $t = 0, T$, the change in p_S/p_G and the growth of total output per worker. For our benchmark analysis we set $\rho = 1.42$, which corresponds to the value used in Katz and Murphy (1992), and which is commonly used in the literature. Though this is a commonly used value in the literature, it is worth noting that previous estimates using an aggregate production function do not necessarily apply in our setting. For this reason we will also do sensitivity analysis with regard to ρ over a fairly wide interval, ranging from 0.77 to 2.50. Table 2 provides values for the inputs into the calibration procedure. Table 3 below shows the implied values for the technology parameters.

Table 2
Values Used to Calibrate Technology Parameters

| f_{L0} | f_{LT} | w_{H0} | w_{HT} | $\% \Delta \frac{p_S}{p_G}$ | $\% \Delta Y$ |
|---------------|---------------|---------------|---------------|-----------------------------|----------------------|
| 0.78 | 0.67 | 1.41 | 1.90 | 62.0 | 70.0 |
| θ_{G0} | θ_{GT} | θ_{S0} | θ_{ST} | $\frac{C_{S0}}{Y_0}$ | $\frac{C_{ST}}{Y_T}$ |
| 0.18 | 0.34 | 0.54 | 0.66 | 0.29 | 0.44 |

Table 3
Calibrated Technology Parameters ($\rho = 1.42$)

| α_{G0} | α_{S0} | α_{GT} | α_{ST} | A_{ST}/A_{S0} | A_{GT}/A_{G0} |
|---------------|---------------|---------------|---------------|-----------------|-----------------|
| 0.28 | 0.55 | 0.43 | 0.66 | 1.25 | 2.27 |

A few remarks are in order. Not surprising given the way in which we grouped industries into the two sectors, we see that the weight on low-skilled labor is greater in the goods sector than in the service sector at both dates. More interesting is that in both sectors technological change has an important component that is skill biased. While the level drop in α is larger for the goods sector than the service sector, the percent changes are very similar, though slightly larger for the service sector than the goods sector (22% versus 20%). It follows that the extent to which technological change is skill biased is relatively similar for the two sectors. However, overall technological progress is much greater in the goods sector than in the service sector. The TFP term in the goods sector more than doubles over the nearly thirty year period from 1977 to 2005, corresponding to an average annual growth rate of 2.97%. In contrast, the growth of the TFP term in the service sector averages only 0.80% per year.

We now turn to the issue of determining the values for the three preference parameters: a_G , \bar{c}_S and ε . While 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. As noted above, the calibration of technology parameters used information about sectoral expenditure shares without guaranteeing that observed expenditure shares were consistent with household demands given all of the prices. Requiring that the aggregate expenditure share for goods (or services) is consistent with the observed values in the data for the initial and terminal date would provide two restrictions on the three preference parameters. It follows that we would either need to introduce an additional moment from the data, or perhaps use information from some outside study to determine one of the three preference parameters. For our benchmark results we will follow the second approach and fix the value of ε , and then use data on aggregate expenditure shares to pin down the values for a_G and \bar{c}_S . It will turn out that our main finding is relatively robust to variation over a large range of values of ε , thereby lessening the need to tightly determine its value. Nonetheless, in Section 5 we will describe how cross-sectional data on expenditure shares could be used as an additional moment and allow us to determine all three preference parameters.

The empirical literature does not provide estimates of ε that correspond to our definitions of the two sectors. However, as noted previously, what we label the goods sectors does contain almost all of the industries that produce goods, while the sector that we label as services does consist primarily of service sector industries in the actual economy. However, this split is not sharp, since low skilled service sector industries such as retail trade are also included in what we label the goods sector. Nonetheless, we believe that information about the elasticity of substitution between the “true” goods and services sectors should be informative about a reasonable 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 (2014). All of these studies suggest very low degrees of substitutability between true goods and true services. For this reason we consider values for ε in the set $\{0.125, 0.20, 0.50\}$, with $\varepsilon = 0.20$ chosen as our benchmark.¹⁷

Simple manipulation of the household demands gives:

$$c_{St} = \frac{\left(\frac{p_{Gt}}{p_{St}}\right)^{\varepsilon-1} \left(\frac{1-a_G}{a_G}\right)^\varepsilon}{\left(\frac{p_{Gt}}{p_{St}}\right)^{\varepsilon-1} \left(\frac{1-a_G}{a_G}\right)^\varepsilon + 1} - \frac{1}{\left(\frac{p_{Gt}}{p_{St}}\right)^{\varepsilon-1} \left(\frac{1-a_G}{a_G}\right)^\varepsilon + 1} \frac{p_{St}\bar{c}_S}{1 - f_t^H + f_t^H w_t}$$

for total demand for services in period t . Using information on relative prices

¹⁷Comin 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, though still less than 0.50, which is our upper range.

and total income, given a value for ε this allows us to determine values for a_G and \bar{c}_S .

Table 4 shows the values for the preference parameters in the different scenarios.

Table 4
Calibrated Preference Parameters

| | ε | a_G | \bar{c}_S |
|--------------------|---------------|-------|-------------|
| Benchmark | 0.20 | 0.72 | 0.14 |
| High ε | 0.50 | 0.31 | 0.35 |
| Low ε | 0.125 | 0.92 | 0.11 |

The qualitative patterns in this table are intuitive. In all cases the change in income, relative prices and the aggregate expenditure shares are all the same. As we move from $\varepsilon = 0.20$ to $\varepsilon = 0.125$ we decrease the elasticity of substitution between the two goods, implying a smaller response in relative quantities but a larger response in relative expenditure shares. In order to compensate for this larger effect, we need to decrease the impact of income changes on relative expenditure shares, implying a lower value for \bar{c}_S . The lower value for \bar{c}_S will in turn lead to a higher expenditure share on services in the initial period, since the non-homotheticity is now less important. Hence, in order to match the expenditure shares for the initial period we need to attach a lower weight to consumption of goods. As we move from $\varepsilon = 0.20$ to $\varepsilon = 0.50$ we see the reverse pattern.

5 Decomposing Changes in the Skill Premium

Having calibrated the model so as to be consistent with growth, structural transformation, changes in the skill premium and changes in the share of skilled labor, we can now use the model to perform counterfactuals that allow us to attribute changes in the skill premium to changes in various aspects of the economic environment. Our primary objective is to decompose the effect of changes in technology on the skill premium into a piece due to skill biased technological change and a residual piece that is due to other forms of technological change. The residual piece affects the relative demand for skilled individuals indirectly, through its impact on the relative consumption of services.

In order to decompose the effects of technological change into these subcomponents we carry out several counterfactuals, which we report in Table 5, for each of the three specifications that differ with respect to the value of ε .

Table 5
Decomposing Changes in the Skill Premium
US, 1977-2005

| | $\varepsilon = 0.50$ | $\varepsilon = 0.20$ | $\varepsilon = 0.125$ |
|---|----------------------|----------------------|-----------------------|
| w_{H0} Data | 1.41 | 1.41 | 1.41 |
| w_{H0} Model | 1.41 | 1.41 | 1.41 |
| w_{HT} Data | 1.90 | 1.90 | 1.90 |
| w_{HT} Model | 1.90 | 1.90 | 1.90 |
| w_{HT} Model—changes in f_i only | 0.96 | 0.92 | 0.92 |
| w_{HT} Model—changes in f_i and A_j only | 1.17 | 1.16 | 1.16 |
| w_{HT} Model—changes in f_i and α_j only | 1.61 | 1.59 | 1.59 |

The first four rows of the table simply report that in the data, the skill premium increased from 1.41 to 1.90 between 1977 and 2005, and that our calibrated model perfectly replicates this change. The rest of the table decomposes this change in the model by considering several counterfactuals. The first counterfactual assesses the role of “supply” versus “demand” factors. Specifically, the share of labor supply coming from skilled workers increases between 1977 and 2005, and in the absence of any other changes exerts downward pressure on the skill premium. The fifth row of Table 5 shows that if the change in relative supply of skill had been the only change over time, it would have resulted in a drop in the skill premium of between 45 and 49 percentage points across the three specifications. Given that we in fact observe an increase in the skill premium of 49 percentage points, it follows that the overall effect of technological change is to increase the skill premium by between 94 and 98 percentage points.

Our next goal is to decompose the overall effect of technological change into the part that is due to skill biased technological change (i.e., changes in the α_{jt} 's) and a residual due to other dimensions of technical change (i.e., the A_{jt} 's). Before doing so it is of interest to note that in standard aggregate analyses of technological change, it is only skill biased technological change that affects the skill premium, since changes in TFP have no impact on relative marginal products holding factor supplies constant. Central to our analysis is the fact that this result does not hold in our model. The reason that this is true is precisely because the two sectors have different skill intensities. As a result, any change in the economic environment that leads to changes in the composition of demand will impact on the relative demand for skill. As is standard in the structural transformation literature, our model features two forces through which changes in the A_{jt} 's can influence the sectoral composition of output. One is through income effects: because of nonhomothetic preferences, any change in technology that increases income will lead to a reallocation of activity from the goods to the services sector, thereby indirectly increasing the relative demand for skill. The other is through relative price effects. If goods and services have an elasticity of substitution that is less than unity, then decreases in the relative productivity of services will also lead to a reallocation of factors of production to the service sector, again indirectly increasing the relative demand for skill.

There are two natural exercises that one could perform to assess the con-

tribution of changes in the A_{jt} 's to changes in the skill premium. In the first we shut down the changes in the α_{jt} 's and compute what fraction of the overall increase is accounted for by changes in the A_{jt} 's alone. In the second we shut down changes in the A_{jt} 's and find the residual that is not accounted for by changes in the α_{jt} 's. In a linear model these two exercises would give the same answer, but to the extent that nonlinearities are present they may differ. It will turn out that the answers do differ, but only to a minor extent. The final two rows in Table 5 present the results of these two counterfactuals. For concreteness, we first focus on the case of $\varepsilon = 0.20$. When we hold the α_{jt} 's fixed, we find that the change in the A_{jt} 's accounts for 24 percentage points of the overall 98 percentage point increase accounted for technical change, or approximately 24%. If we instead fix the A_{jt} 's, the residual is 31 percentage points, which represents approximately 31% of the total. Based on this we conclude that non-skill biased technical change accounts for between 25 and 30 percent of the overall change in the skill premium due to technical 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 18 percentage points instead over the period 1977 to 2005 instead of increasing by 49 percentage points.

If we redo this calculation for the other two values of ε the answers are similar. For $\varepsilon = 0.50$ the two methods imply that changes in the A_{jt} 's account for 22% and 31% of the overall change in the skill premium due to technical change, whereas for $\varepsilon = 0.125$ the two values are effectively identical to those for the $\varepsilon = 0.20$ case, being equal to 24% and 31%. From this we conclude that our finding of significant contribution of changes in the A_{jt} 's is robust to a large variation in the value of ε .

In the introduction we stressed the fact that aggregate production function analyses abstract from compositional changes, and that our main objective was to assess the quantitative importance of the compositional changes that are associated with the process of structural transformation during development. In order to illustrate that this is the mechanism that we are picking up in the above calculations it is of interest to examine the relationship between the different changes in the economic environment and the sectoral composition of value added. Table 6 reports the results for each of the three values of ε .

Table 6
 Technical Change and Value Added Share of Services
 US, 1977-2005

| | $\varepsilon = 0.50$ | $\varepsilon = 0.20$ | $\varepsilon = 0.125$ |
|-----------------------------|----------------------|----------------------|-----------------------|
| US Data 1977 | 0.29 | 0.29 | 0.29 |
| Model 1977 | 0.29 | 0.29 | 0.29 |
| US Data 2005 | 0.44 | 0.44 | 0.44 |
| Model 2005 | 0.44 | 0.44 | 0.44 |
| Model with fixed A_j | 0.26 | 0.25 | 0.25 |
| Model with fixed α_j | 0.44 | 0.46 | 0.46 |

The first four rows of the table simply confirm that the service sector grew significantly between 1977 and 2005, increasing its share of value added from 29 percent to 44 percent, and that our calibrated model perfectly accounts for this change. The last two rows provide two different ways of assessing the role of changes in the A_j 's and the α_j 's in accounting for this compositional change. Both methods provide the same simple message: virtually all of the compositional change is accounted for by changes in the sectoral TFPs. It follows that our previous decomposition of changes in the skill premium due to the two different sources of technical change can effectively be interpreted as statements about the importance of compositional changes.

Non-skill biased technological change in our model still has two dimensions: one of which increases the overall level of TFP in the economy and the other of which led to higher relative TFP in the goods sector. As we noted above, both of these changes tend to reallocate activity from the goods sector to the service sector, thereby indirectly increasing the relative demand for skill. It is perhaps of interest to examine the relative magnitude of these two effects. Note that for given changes in the A_{jt} 's the relative magnitude of these two effects is dictated by the preference parameters ε and \bar{c}_s : as ε becomes smaller, relative TFP changes have larger effects, and as \bar{c}_s becomes larger then sector neutral changes in the A_{jt} 's have larger effects. Because our calibration procedure implies that as ε becomes smaller the value of \bar{c}_s decreases, we expect to find that sector neutral change plays a larger role for smaller values of ε . To evaluate this we consider the counterfactual in which we hold all parameters fixed from the original calibration, allow the f_{it} 's and the α_{jt} 's to change as before, but counterfactually force the A_{jt} 's to grow at the same rate so as to yield the same overall change in aggregate output as in the data. When we do this, the implied values of the skill premium are 1.83, 1.72, and 1.69 for the cases of $\varepsilon = 0.50, 0.20,$ and 0.125 respectively. It follows that when $\varepsilon = 0.50$ it is income effects that dominate the overall impact of the A_{jt} 's on the skill premium, whereas for the smaller values of ε the sector biased nature of TFP growth is somewhat more important than the income effect. So while the three different specifications offer very similar decompositions regarding the overall effect of changes in the A_{jt} 's, they do have different implications for what type of changes in the A_{jt} will lead to future changes in the skill premium.

5.1 Using Data to Infer ε

In the results above we considered a range of values for ε rather than trying to use our model to infer a specific value. Since our main message was robust to a wide range of values for ε we do not view this as a particular limitation of the analysis. Nonetheless, in this subsection we describe how the use of a cross-sectional moment on the household side would allow us to also infer a value for ε . One way to frame the issue is the following. The expenditure shares at time 0 effectively pins down the weight on goods in the utility function. Our calibration procedure requires that changes in aggregate expenditure shares be consistent with observed changes in income and relative prices. However, all this

does is restrict us to a one parameter family of combinations of income effects and substitution effects, i.e., there is a one parameter family of pairs (ε, \bar{c}_S) that can deliver the required change in expenditure shares for given changes in income and relative prices. In order to uniquely pin down the parameter pair we would need an additional moment that reveals information about the size of these two effects.

Intuitively, cross-sectional information can provide information about the magnitude of the income effect: assuming that all households have the same preferences and face the same prices at a point in time, cross-sectional heterogeneity in income will allow us to infer the size of the income effect. This can be implemented using the cross-sectional information that we reported in Section 2. In particular, our empirical analysis found that the expenditure share for services is 0.04 higher for high skill households than for low skill households. If we require that our model match this moment in the final time period, the implied values for ε and \bar{c}_S are 0.05 and 0.08 respectively when we assume $\rho = 1.42$. This would correspond to values of ε somewhat below the lower end of the interval that we considered. As a practical matter, it turns out that moving to increasingly smaller values of ε from $\varepsilon = 0.20$ has relatively small effects, as could already be seen in Table 4. We have also repeated this exercise for alternative values of ρ and obtained estimates of $\varepsilon = 0.11$ and 0.01 for ρ equal to 0.77 and 2.50 respectively. While we do not report the results for these cases in detail, we note that our main message remains unaffected if we were to adopt these specifications.

Having offered the idea of using cross-sectional data to infer the size of income effects on the demand for what we have labelled goods and services, we think it is important to repeat one important limitation of this approach in the current context. Two of the largest components of high skill services are education and health care, both of which are not well tracked by household expenditure surveys. To the extent that spending on some components of education and health reflect a collective societal choice, it is not clear that cross-sectional data will be useful in detecting how cross-sectional differences in income affect desired consumption.

5.2 Sensitivity Exercises

For the results in the previous section we assumed that $\rho = 1.42$, which we noted was a standard value in the literature, and the value implied by the analysis in Katz and Murphy (1992). However, we also noted that the aggregate analyses that have supported this estimate are not necessarily appropriate in our multi-sector economy. For this reason we also consider a wider range of values for ρ to assess the extent to which the above conclusions are robust to variation in this parameter.

We consider two alternative values of ρ , corresponding to higher and lower elasticities of substitution. Specifically, we consider $\rho = 0.77$ and $\rho = 2.5$. In each case we redo the calibration procedure as before. While the value of ρ does affect the quantitative findings, it leaves our main message largely unchanged.

For example, focusing on the case of $\varepsilon = 0.20$ we find that when $\rho = 0.77$, the share of changes in the skill premium due to technical change that are accounted for by changes in the A_{jt} is 27% and 41% from the two methods. When $\rho = 2.50$ the corresponding values are 19% and 23%. We conclude that our main finding of a significant role for changes in demand composition induced by technical change in accounting for changes in the skill premium is robust to considering a wide range of values for ρ , though higher values of this elasticity parameter do lead to modest declines in the estimated role played by demand composition.

To this point our analysis has implicitly assumed that the value of ρ is the same in both sectors, which given the absence of any empirical evidence on the extent of heterogeneity in ρ across sectors, seemed a natural benchmark. Reshef (2013) suggests that the elasticity of substitution between high and low-skilled workers may be lower in services, for example. It is therefore important to assess whether our results are sensitive to the equality assumption. To do this we redo our exercise for several specifications in which we allow the two values of ρ to vary across sectors, and allowing for the ratio ρ_G/ρ_S to be both larger and smaller than one. In all cases we assume that the weighted average of the two elasticities— $(H_G/H)\rho_G + (H_S/H)\rho_S$ —is equal to 1.42 when evaluated at the initial factor shares, so that our analysis can be interpreted as assessing the effect of heterogeneity holding the aggregate elasticity of substitution constant. We consider values for ρ_S of 0.77, 0.91, 1.11, and 2.00, and the implied values for ρ_G are 2.23, 2.06, 1.82, and .73. Table 7 reports the same statistics as in Table 5, focusing on the case of $\varepsilon = 0.20$.

Table 7
The Effect of Sectoral Variation in ρ
US, 1977-2005

| | $\frac{\rho_S}{\rho_G} = 0.35$ | $\frac{\rho_S}{\rho_G} = 0.44$ | $\frac{\rho_S}{\rho_G} = 0.61$ | $\frac{\rho_S}{\rho_G} = 1.00$ | $\frac{\rho_S}{\rho_G} = 2.73$ |
|---------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|
| w_{H0} Data | 1.41 | 1.41 | 1.41 | 1.41 | 1.41 |
| w_{H0} Model | 1.41 | 1.41 | 1.41 | 1.41 | 1.41 |
| w_{HT} Data | 1.90 | 1.90 | 1.90 | 1.90 | 1.90 |
| w_{HT} Model | 1.90 | 1.90 | 1.90 | 1.90 | 1.90 |
| Counterfactual w_{HT} | | | | | |
| changes in: f_i only | 1.00 | 0.99 | 0.97 | 0.92 | 0.86 |
| f_i and A_j only | 1.15 | 1.15 | 1.15 | 1.16 | 1.18 |
| f_i and α_j only | 1.66 | 1.66 | 1.64 | 1.59 | 1.46 |

For ease of comparison, the fourth column repeats the results from our benchmark specification. For values of $\rho_S/\rho_G < 1$ the implications are affected very little, and to the extent that a very large value of ρ_S/ρ_G influences the quantitative results, it yields a larger role for the demand side effects that we focus on. Noting that we are considering a very wide range of variation in the relative values of ρ , we conclude that our results are quite robust to variation in ρ across sectors.

Lastly, 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 high skill intensive sector. Between 1977 and 2005 this relative price increased by more than sixty percent. One possible concern is that the price inflation in the high skill intensive sector might be upward biased because of the failure to properly account for quality improvements. 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 high 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 15 additional percentage points. We repeat our benchmark exercise for the case of $\rho = 1.42$ and $\varepsilon = 0.20$, carrying 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 demand factors 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.

6 Decomposing Changes in Relative Prices

While our main focus has been to use our model to understand the relative importance of different factors in generating the observed changes in the skill premium, our model also allows us to assess the importance of different factors in generating the change in the relative price of skill intensive services over time. In particular, our model suggests two distinct channels at work. As is standard in the literature on structural change with uneven technological progress across sectors, differential growth in sectoral TFP will lead to changes in relative sectoral prices. But our model also features an additional channel: because the sectors have different factor shares, changes in the relative price of factors will also lead to changes in relative sectoral prices. In particular, since the high-skill intensive sector uses skilled labor more intensively, any increase in the relative price of skilled labor will lead to a higher relative price for this sector. This effect was previously documented in equation (8).¹⁸

Here we consider some counterfactuals for our benchmark specification in which $\rho = 1.42$ and $\varepsilon = 0.20$ in order to assess the relative importance of the different forces. In the data, the relative price of high-skill intensive services increases by 62 percentage points between 1977 and 2005, and by virtue of our calibration procedure, our model perfectly accounts for this increase. To assess the pure role played by the increase in the skill premium, we compute the implied relative price from equation (8) assuming that all technology parameters remain fixed at their 1977 values, but letting the skill premium increase from

¹⁸Buera and Kaboski (2012) highlight this effect in a theoretical model in which the difference in skill-intensity across the goods and service sectors arises endogenously.

1.41 to 1.90, as in the data. The result is an increase in the relative price of skill intensive services of 11 percentage points, or roughly 18% of the overall increase.

It follows that the direct effects of technological change are the dominant force behind the increase in the relative price of skill intensive services. Moreover, it is the differences in sectoral TFP growth that drives this direct effect. Again, if we take equation (8) and now instead hold the skill premium and sectoral TFP constant and consider the pure effect of skill biased technological change, the result is that the relative price of skill intensive services would have *decreased* by 19 percentage points.

This last calculation examined the direct effect of changes in skill-biased technological change, but without incorporating the general equilibrium effect on wages. Our previous counterfactuals argued that the contribution of skill biased technological change to the skill premium served to increase it from 1.41 to 1.59. If we include this effect in combination with the direct effect of skill biased technological change, the result is that the relative price of skill intensive services decreases by 15 percentage points.

In summary, we conclude that although increases in the skill premium directly account for a non-trivial share of the increase in the relative price of high-skill intensive services, the dominant factor behind this increase is the relatively slow sectoral TFP growth in this sector.

7 Cross Country Analysis

In this section we use data from nine other OECD countries for which the available data exists to address two distinct issues. The first issue concerns model validation, and the second issue is to assess the importance of skill-biased structural change for a larger set of countries.

7.1 Model Validation Using Cross-Country Data

Our calibration procedure assigned parameter values by targeting the same number of moments as there were parameters. While both the production structure and our method for inferring technological change are very standard, taking the production side as given we assumed that the values for the preference parameters of our utility function were such that the model would match the beginning and end points of the data for valued added expenditure shares. Even if our utility function were mis-specified in an important way, this procedure would still allow us to fit the beginning and end points, but in this case we might be wary of using our calibrated specification for the counterfactual exercises. One simple test of the specification is to consider its ability to fit not only the two endpoints of our sample, but also the actual time series. Unfortunately this is not a very stringent test for the period we are studying, since the key series in our analysis are fairly linear, and the model is able to match them fairly well.

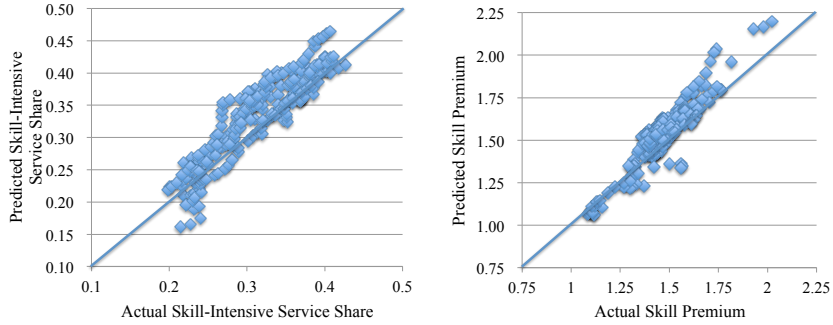


Figure 3: Model Fit in a Panel of Countries: Structural Change (left panel) and the Skill Premium (right panel).

As a somewhat more stringent test, we turn to cross country data. For this exercise we use data from the following nine countries: Australia, Austria, Denmark, Germany, Italy, Japan, the Netherlands, Spain and the United Kingdom. We assume that the utility function for each country is the same as the one implied by our benchmark calibration with $\rho = 1.42$ and $\varepsilon = 0.20$, i.e., we impose the implied values for a_G and \bar{c}_S . Additionally, we assume that ρ is the same for all countries. However, using the same procedure as above, for each country we use our model to infer the time series for technological change, as well as measuring the supply of skill in each country. Because preference parameters are imported from the calibration using US data, we have not imposed that the model will fit the time series of interest for each country. Nonetheless, Figure 3 shows that this specification provides a reasonably good fit to the actual data for this set of countries. Because the behavior of the skilled labor share and the skill premium do differ across countries, we believe that this finding is supportive of our parsimonious structure.¹⁹

It is also of interest to note that the above procedure also implies processes for technological change that are broadly similar across countries over development, as shown in Figure 4.²⁰ To the extent that we believe the process of technology adoption and diffusion are at least generally similar across richer countries, we would view it as somewhat problematic if our procedure indicated dramatically different processes across these countries.

¹⁹We note that our model does not do such a good job of matching the series for Korea. Notably, it specifically fails for the early part of the sample in which Korea has very low GDP. We interpret this as evidence that our specification of non-homotheticities is probably best viewed as an approximation that holds in a restricted range of incomes, and that if one wants to consider a much larger range of incomes then it is probably important to consider more general specifications such as those in Boppart (2014) and Comin et al. (2015).

²⁰The plots in Figure 4 have removed country fixed effects in order to focus on the changes in technology over time rather than the cross-sectional differences.




Figure 4: Calibrated Technological Processes: Sector-Biased (left panel) and Skill-Biased (right panel) Technologies. The diamonds (squares) correspond to the high (low) -skill intensive sector.

7.2 Skill-Biased Structural Change and the Skill Premium in Cross Country Data

In this subsection we assess the extent to which skill biased structural transformation has influenced the skill premium in each of the countries that we studied in the previous subsection. We could carry out this calculation for the specifications in the last subsection, i.e., assuming the same preference parameters for these countries as in the US. A potential disadvantage of this method is that although the model with common preference parameters across countries offers a good fit to the cross country time series data, it does not necessarily account for all of the changes in the skill premium for each of the countries. Alternatively, we could assume country specific values for a_G and \bar{c}_S and simply repeat the analysis that we have carried out for the US for each of the additional economies. These two methods provide fairly similar answers, and in the interest of space we only report here the results of the second exercise, which are shown in Table 8. To compute the values in Table 8 we first calculate the contribution of all forms of technological change by computing the difference between the actual skill premium in 2005 versus the skill premium that would have existed in 2005 if there had been no technological change relative to 1977 but allowing for the observed change in the supply of skill. We then isolate the fraction of this overall contribution of technological change that is due to skill biased structural change by computing the fraction of this change that is accounted for by changes in the A_j 's.

Table 8

| Contribution of SBTC Across Countries | |
|---------------------------------------|------|
| Australia | 0.18 |
| Austria | 0.40 |
| Denmark | 0.11 |
| Spain | 0.32 |
| Germany | 0.37 |
| Italy | 0.54 |
| Japan | 0.22 |
| Netherlands | 0.27 |
| UK | 0.36 |

The magnitude of this share varies significantly, from a low of 11% in Denmark to a high of 54% in Italy, but the mean value of 31% is very much in line with our estimates from the US. We conclude that the demand side forces associated with skill biased structural change seem to be quantitatively significant in a broad group of advanced economies.

8 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 high-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 technological change overall increased the skill premium by almost 100 percentage points, and that between 25 and 30 percent of this change is due to technological change which operated through compositional changes.

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 high-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. As we noted in Section 2, 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. The literature has also emphasized the possibility that increases in trade might lead to changes in the composition of valued added across sectors. Katz and Murphy (1992) specifically noted this possibility, and more recent analyses include Feenstra and Hanson (1999) and Autor et al. (2013). We think it is important to note that the compositional effects we have focused on are not likely to be reflecting changes due to trade. The reason

for this is that our high-skill intensive sector is composed entirely of industries from the service sector. It is plausible that part of what we identified as within sector skill biased technical change may at least in part reflect compositional effects due to trade, to the extent that trade had caused manufacturing activity in the US to shift to more skill intensive industries.

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.
- AUTOR, D. H. AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103, 1553–97.
- 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.
- 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.
- (2015): “To Which Extent is the Rise in the Skill Premium Explained by an Income Effect?” Working paper, IIES Stockholm University.
- 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.
- COMIN, D., D. LAKSHARI, AND M. MESTIERI (2015): “Structural Change with Long Run Income and Price Effects,” Manuscript.
- DI NARDO, 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.
- 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. RÍOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68, 1029–1054.
- LEONARDI, M. (forthcoming): “The Effect of Product Demand on Inequality: Evidence from the US and the UK,” *American Economic Journal: Applied*.
- 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.
- PIERCE, B. (2010): “Recent Trends in Compensation Inequality,” in *Labor in the New Economy*, National Bureau of Economic Research, Inc, 63–98.
- RESHEF, A. (2013): “Is Technological Change Biased Towards the Unskilled in Services? An Empirical Investigation,” *Review of Economic Dynamics*, 16, 312–331.
- SWIECKI, T. (2014): “Determinants of Structural Change,” Manuscript, University of British Columbia.