Trading Up and the Skill Premium*

Nir Jaimovich[†], Sergio Rebelo[‡],
Arlene Wong,[§] and Miao Ben Zhang[¶]

May 2019

Abstract

We study the impact on the skill premium of increases in the quality of goods consumed by households ("trading up"). Our empirical work shows that high-quality goods are more intensive in skilled labor than low-quality goods and that household spending on high-quality goods rises with income. We propose a model consistent with these facts. This model accounts for the past rise in the skill premium with more plausible rates of skill-biased technical change than those required by the canonical model. It also implies that an expansion of the skilled labor force reduces the skill premium by much less than in the canonical model.

J.E.L. Classification: J2, O4.
Keywords: Quality, skill premium, growth, inequality.

---

*We thank Daron Acemoglu, Gadi Barlevy, Carlos Burga, Martin Eichenbaum, Berthold Herrendorf, Eric Hurst, Ben Moll, Jonathan Vogel, and Carlo Zanella for their comments.

[†]University of Zurich and CEPR.
[‡]Northwestern University, NBER, and CEPR.
[§]Princeton University and NBER.
[¶]University of Southern California.

1 Introduction

U.S. income inequality has greatly increased in the last four decades. This increase has motivated a plethora of policy proposals aimed at narrowing the gap between rich and poor. These proposals include making income taxes more progressive (Diamond and Saez (2001) and Landais, Picketty and Saez (2011)), introducing wealth taxes (Saez and Zuckman (2019)), subsidizing college tuition for students from low-income households (Chetty et al. (2017)), and investing in neighborhoods to promote upward mobility (Chetty and Hendrem (2018)).

In evaluating these and other policy proposals, it is useful to understand the dynamics of income inequality: are there stabilizing forces that naturally narrow the gap between rich and poor? One such force is the likely increase in the relative supply of high-skill workers. In the canonical model used by Katz and Murphy (1992) and many others, this increase lowers the skill premium, naturally reducing income inequality.

In this paper, we argue that this stabilizing force is likely to be weaker than suggested by the canonical model. This weakness stems from the fact that as income rises households "trade up," that is, they increase the quality of the goods and services they consume. Trading up fosters inequality in labor earnings because, as we show in this paper, high-quality goods are more intensive in skilled labor than low-quality goods. So, as households improve the quality of what they consume, they increase the demand for skilled labor, contributing to a rise in the skill premium.

This paper has an empirical and a theoretical component. In our empirical work, we show that skill intensity rises with quality and that the quality of household consumption rises with income.

To study the relation between quality and skill intensity, we proceed as follows. We use the price of a good or service as a proxy for its quality. The idea underlying this approach is that if consumers are willing to pay more for an item, they perceive it to be of higher quality. Our price measures come from two sources: Nielsen Homescan and

Yelp!, a website where consumers post review information about different goods and services.

We match each establishment in the Yelp! data and each manufacturing firm in the Nielsen Homescan data with the Microdata of Occupational Employment Statistics (OES) from the Bureau of Labor Statistics. We combine these data with the U.S. Department of Commerce Current Population Survey (CPS) to construct measures of skill intensity for each establishment and manufacturing firm in our data.

To document the relation between quality and income, we use data from Nielsen Homescan as well as data on durable expenditures from the Consumer Expenditure Survey (CEX).

In our theoretical work, we propose a simple model consistent with our two facts. The model is designed to be as similar as possible to the canonical model used by Katz and Murphy (1992) to study the dynamics of the skill premium. In our model, both skill-biased technical change (SBTC) and Hicks-neutral technical change (HNTC) increase real income, thus expanding the demand for quality. Since quality is intensive in high-skill workers, the demand for these workers rises, leading to an endogenous rise in the skill premium and a widening of the income distribution.

Our model has implications for how we interpret the past dynamics of the skill premium. Using Fernald's (2014) estimates of the rate of HNTC, we compute the rate of SBTC consistent with the rate of change in the quality of goods consumed estimated by Bils and Klenow (2001).[1] We find that our model accounts for the observed rise in the skill premium in the last four decades with an annual rate of SBTC of 1.05 percent per year. In contrast, the canonical model used by Katz and Murphy (1992) requires a rate of SBTC of 5.5 percent per year.

Our model also has implications for the future dynamics of the skill premium. To explore these implications, we forecast the increase in the supply of skilled workers.

[1]These estimates are obtained by regressing the Bureau of Labor Statistics' inflation measure on the growth rate of expenditures instrumented with the slope of the quality Engel curve.

To isolate the effect of this supply increase, we abstract from HNTC and SBTC. The canonical model implies a 21 percent fall in the skill premium. In contrast, our model predicts a 15 percent fall in the skill premium. This smaller decline in the skill premium occurs because the preponderance of high-skill workers in the labor force is associated with higher household income which boosts the demand for quality, increasing the demand for skilled labor.

To explore these results and study the robustness of our conclusions, we consider two variants of the benchmark model. The first variant has two types of households, one composed by high-skill workers and the other by low-skill workers. These households consume different levels of quality, so there are two quality levels produced in equilibrium. We find that the results are similar to those of the benchmark model. The reason for this similarity is that in our basic calibration the price-quality schedule is not far from linear, so producing two levels of quality instead of one does not affect substantively the properties of the equilibrium.

In the second variant of the benchmark model, the household consumes two types of goods, a basic good with a constant level of quality and a quality good, whose quality can vary on a continuum. The level of SBTC required to explain the rise in the skill premium in this model is in between that implied by our benchmark model and the canonical model, but closer to the benchmark model. These properties result from the presence of two effects. First, a rise in income is spent in part on the basic good, weakening the degree of trading up. Second, since the quality good is superior, its share in spending rises with income. This effect increases the impact of income rises on the degree of trading up and on the skill premium.

Our work is related to five strands of literature. The first strand is the large body of work on SBTC surveyed by Acemoglu and Autor (2011). The second strand is the literature on the role of capital deepening in explaining the evolution of the skill premium (e.g. Krusell, Ohanian, Ríos-Rull, and Violante (2000), Polgreen and Silos (2008), and Burstein and Vogel (2017)). We show in the Appendix that incorporating the trading-up phenomenon in these models reduces the resulting estimates of the elasticity of substitution between unskilled labor and capital. The third strand is work on skill-biased structural change (e.g. Acemoglu and Guerrieri (2008), Buera and Kaboski (2012), Buera, Kaboski and Rogerson (2015), Boppart (2015), and Alon (2018)). This work emphasizes how rises in income shift demand towards sectors that are more intensive in skilled work. In contrast, we emphasize that, as income rises, the demand for quality increases, raising the demand for skilled labor within a given sector. One important difference between these two mechanisms is that the process of upgrading quality within a sector is presumably unbounded while sectoral reallocation is likely to be bounded. The fourth strand is work on the importance of quality choice in growth models (e.g. Grossman and Helpman (1991a,b), Stokey (1991), Zweimüller (2000), and Foellmi and Zweimüller (2006)), trade models (e.g. Verhoogen (2008) and Fieler, Eslava and Xu (2017)) and macro models (e.g. Jaimovich, Rebelo, and Wong (2019)). The fifth strand is the industrial-organization literature on product differentiation (e.g. Shaked and Sutton (1982) and Berry (1994)).

This paper is organized as follows. Sections 2 and 3 contain our empirical work. In Section 4 we present our benchmark model and discuss its implications for the rate of SBTC required to explain the rise in the skill premium. We also consider an extension of the model with two heterogeneous households. In Section 5 we analyze the model with two goods, one with constant quality and the other with variable quality. Section 6 concludes.

2 Quality and Skill Intensity

In this section, we measure the intensity of skilled labor in establishments that are in the same sector but produce products of different quality. To accomplish this goal, we first construct measures of both the quality of goods produced and skill intensity by establishment. We then study the relation between quality and skill intensity.

There are three approaches used in the literature to measure quality. The first approach, which we adopt in this paper, is to use relative prices as proxies for quality. The idea underlying this approach is that if consumers are willing to pay more for an item, they perceive it to be of higher quality.[2] The second approach, is to infer quality from the materials and labor costs used in production (e.g. Veerhoogen (2008)). The third approach, is to structurally estimate quality using data on prices and quantities combined with functional form assumptions about household utility (e.g. Veerhoogen (2008) and Hottman, Redding and Weinstein (2016)). We focus on relative prices as measures of quality because it allows us to use a broader sample of goods and of firms included in the OES data set.[3]

There is strong evidence that relative prices are positively correlated with the quality measures produced by the other two approaches. For example, Veerhoogen (2008) shows that higher-quality items, which have higher costs of production, also have higher prices. Hottman, Redding and Weinstein (2016) and Khandelwal (2010) find that quality is strongly positively correlated with relative prices within product groups.

Before we dive into our empirical analysis, it is useful to illustrate our results with data for the restaurant sector. According to OES data, the key occupations in restaurants are: managers and executives, chefs and head cooks, first-line supervisors of food preparation, cooks and food preparation workers, waiters and waitresses, serving workers, and marketing and sales. Chefs, a high-skill occupation, represents on average 20 percent of the workforce in a full-service restaurant but only 2 percent in a limited-service restaurant.[4] As income rises and consumers trade up in the quality of the restaurants they patronize, the demand for chefs expands, contributing to a rise in the skill premium.

[2]See Jaimovich, Rebelo and Wong (2019) for evidence that supports this assumption.

[3]Structural estimation generally requires price shifters to instrument for price changes. A commonly used price shifter for an item sold in a particular county is the average price of the item in other counties. Using this shifter restricts the analysis to items sold in many counties which results in a substantial reduction in sample size.

[4]At Alinea, a high-end restaurant in Chicago, chefs represent 30 percent of the workforce.

2.1 Measuring Skill Intensity

Our measure of the intensity of skilled labor is based on two data sets. The first is the OES data collected by the Bureau of Labor Statistics (BLS). It covers 1.1 million establishments, representing 62 percent of total employment and spanning all sectors of the North American Industry Classification System at a 6-digit level. Unfortunately, the OES does not contain information on worker education attainment. For this reason, we calculate the distribution of employees across twelve wage bins for each occupation and establishment from the OES and relate this wage distribution to the wages of skilled workers estimated using the CPS. We use the information regarding wages, education and industry in the CPS as follows. For every industry in the CPS, we compute the average wage of college graduates. We classify a worker in the OES as skilled if her wage exceeds the average wage of college graduates for her industry. We then compute for each establishment in the OES data the fraction of employment and of the wage bill accounted for by skilled workers.

We proceed similarly to construct two other measures of the share of skilled workers. For our second measure, we classify workers as skilled if their wage exceeds the average wage of workers with "some college or more" for the worker's industry. For our third measure, we classify workers as skilled if their wage exceeds the average wage for all workers in the respective industry.

Table 1: Establishments' Share of Skilled Workers

Sample	#Est.	Skilled 1		Skilled 2		Skilled 3	
		Emp	Wage	Emp	Wage	Emp	Wage
All Sectors	1,131,170	16.7	36.9	23.7	45.6	27.7	49.9
NAICS Sector:							
Management	13,997	50.3	53.6	63.5	59.5	61.0	63.0
Educational	39,385	33.6	25.4	38.0	38.2	40.9	48.0
Information	33,176	29.3	45.4	34.8	58.2	40.0	64.3
Utilities	6,217	29.8	30.3	35.9	31.1	55.9	31.6
Professional	106,407	28.9	29.1	34.3	38.1	37.6	48.6
Transportation	43,934	28.3	25.7	40.5	37.1	46.6	42.0
Construction	82,188	23.8	44.2	29.4	51.7	38.7	71.7
Finance	56,599	23.6	53.8	30.1	59.6	31.9	64.9
Wholesale	86,176	19.8	31.3	26.8	41.6	30.9	51.6
Health Care	124,463	16.4	55.1	27.1	59.8	29.7	63.0
Other Services	73,062	15.4	10.4	24.4	22.7	27.8	24.7
Manufacturing	107,826	13.9	43.1	20.9	49.4	29.8	59.6
Entertainment	26,549	12.0	38.9	20.0	53.2	19.7	55.5
Real Estate Rental	37,750	10.3	49.9	16.1	56.8	24.8	58.7
Retail	121,065	9.6	42.7	17.8	52.1	21.7	56.1
Administrative	77,873	8.8	74.6	17.2	84.1	25.6	82.3
Accommodation	50,700	3.2	31.7	10.4	43.4	11.5	43.3

Notes: This table shows the average share of skilled workers per establishment by sector computed using OES surveys of establishments from May 2004 to November 2007. We consider three definitions of skilled workers. In "Skill 1, 2 and 3" a worker is skilled if his/her wage exceeds the average wage of college graduates in the industry, the average wage of workers with some college education in the industry, and the average wage of workers in the industry, respectively. The three average wages for each industry are estimated with CPS data. Columns 3, 5 and 7 (labeled "Emp") refer to the average employment share of skilled workers per establishment. Columns 4, 6 and 8 (labeled "Wages") refers to the average share of wages received by skilled workers per establishment. See text for more details.

Table 1 displays our results for different sectors. Consider for example the sector of manufacturing. Using the first measure of skill, we find that the fraction of manufacturing workers who are skilled is 13.9 percent and that these workers earn 43.1 percent of the wage bill. Using the second, broader classification of skill, we find that the fraction of manufacturing workers share of the wage bill earned by these workers are 20.9 and 49.4 percent, respectively. Using the third and broadest classification of skill, we find that the fraction of manufacturing workers share of the wage bill earned by these workers are 29.8 and 59.6 percent, respectively.

2.2 Quality Measures

We use prices of goods as a proxy for their quality. Our price measures come from two sources. The first is data from Yelp!, a website where consumers post review information about different goods and services. The second, is Nielsen Homescan data.

2.2.1 Yelp!-based Quality Measures

For each store and location pair, Yelp! asks users to classify the price of the goods and services they purchased into one of four categories: $ (low), $$ (middle), $$$ (high), and $$$$ (very high). Because there are few observations in the very-high category, we merge the last two categories into a single high-price category.

We match Yelp! establishments to the OES establishments in three steps. First, we match the contact phone numbers of Yelp! establishments to the contact phone number of the OES establishments. Whenever this method does not work, we base the match on name, industry (NAICS 3-digit), and zip code.[5] When multiple OES establishments are matched to one (or multiple) Yelp! establishment(s), we average the skill measures

[5]A major challenge is that the phone number in the OES database can be either the phone number of the establishment or the phone number of the contact person for the establishment, such as the person's mobile phone number. We obtain industry codes of Yelp establishments by matching the Yelp establishments to the ReferenceUSA database, which covers the near universe of establishments in the U.S. Establishments from the two data sets with the same industry, zip code, and similar names based on bigram are also matched. This fuzzy name matching increases the match rate of the two data sets.

of all the OES establishments that are matched to each Yelp! establishment, and assign that average skill measure to the Yelp! establishment.[6] Third, because zip codes are not available for every establishment in the OES database, we conduct a matching procedure similar to that used in our second step based on name, industry (NAICS 3-digit) and county. With these three steps, we obtain the share of skilled labor for 9,908 Yelp! establishments. These data covers the retail, accommodation, entertainment, and information services sectors.

2.2.2 Nielsen-based Quality Measures

In order to extend our analysis to the manufacturing sector, we use the Nielsen Homescan data. This data set contains prices paid and quantities of groceries purchased at a barcode (UPC) level for 113 thousand households over the period 2004-10. Nielsen organizes bar codes into 613 product modules according to where they would likely be stocked in a store.

To construct a measure of quality for each manufacturing firm, we proceed as follows. First, we link each item k (defined at a UPC level) in Nielsen with the manufacturing firm f that produced the UPC using information obtained from GS1 US.[7] We focus on the 2006 data set to match the sample period of the OES data.

Second, we compute the sales-weighted average price across all transactions made during the month t, p_{kft}, for each item k produced by firm f. Similarly, for each product module $m(k)$ that item k belongs to, we calculate $p_{m(k)t}$, the sales-weighted average price within the product module. For each item k, we then calculate the price p_{kft} in month t relative to the average price in the product module, $p_{m(k)t}$:

$$R_{kft} = p_{kft}/p_{m(k)t}.$$

[6]For instance, two Starbucks coffee shops may locate in the same zip code. In this case, our approach assumes that the two coffee shops share the same skill measure.

[7]We thank David Argente, Munseob Lee and Sara Moreira for sharing their code to link UPCs to firms with us. These links between items and firms are also used in Hottman, Redding and Weinstein (2016) and Argente, Lee and Moreira (2019).

By dividing prices by the average price in the product module, we can compare the relative prices of items across different categories of goods.

For single-product firms, our measure of quality is the average relative price for each given item produced by the firm in 2006. For multi-product firms, we compute the firm f's relative price as a weighted average of the relative price of different products, weighted by sales in 2006 (w_{kf}):

$$R_{f,2006} = \sum_{k \in \Omega} w_{k,f,2006} R_{k,f,2006},$$

where Ω denotes the set of all products in the Nielsen data.

Third, we link each manufacturing firm f in the Nielsen-GS1 database to the OES firms. To do so, we perform a fuzzy merge of the first component of the firm names from GS1-Nielsen firm with the first component of the OES legal or trade names. We take a conservative approach in classifying a successful fuzzy merge. A merge between Nielsen-GS1 and the OES is defined as successful if one of the following three situations occur: (i) the similarity score is above 95 percent, (ii) the similarity score, computed using the Levenshtein distance, is above 90 percent and two names share the same first two words; or (iii) the similarity score is above 85 percent and one name contains the other name. This approach yields about 1,600 firms, which include over 29,000 OES establishments.

2.3 Quality and Skill Intensity

Table 2 documents our first fact using the Yelp! data set: the share of high-skill workers employed increases with the quality of the goods produced by the firm. Consider, for example, the results we obtain using our measure of skilled workers based on the average wage of college graduates in the industry. The fraction of high-skill workers is 3.54, 6.38 and 9.49 in our low-, middle- and high-quality tier, respectively. Our estimates of skill intensity naturally vary with the breath of the definition of skill. But, as Table 2 shows, our three alternative definitions of high skill generate similar estimates of the differences in skill intensity of high versus low quality goods.

Table 2: Quality and share of high-skill workers in retail

Sample	#Est.	Skilled 1		Skilled 2		Skilled 3	
		Emp	Wage	Emp	Wage	Emp	Wage
Yelp Sample	9,908	6.01	16.9	13.94	29.02	15.40	31.14
By Quality:							
Low	2,316	3.54	11.15	9.60	21.32	11.48	23.81
Middle	6,089	6.38	17.28	14.94	30.19	16.01	31.80
High	1,503	9.49	23.72	19.40	36.97	21.53	40.24

Notes: This table shows the average share of skilled workers per establishment per Yelp! quality tier. Low quality refers to the $ Yelp! category. Middle quality refers to the $$ Yelp! category. High quality refers to the $$$ and $$$$ Yelp! categories. The statistics are based on the OES survey of establishments from May 2004 to November 2007. We consider three definitions of skilled workers. In "Skill 1, 2 and 3" a worker is skilled if his/her wage exceeds the average wage of college graduates in the industry, the average wage of workers with some college education in the industry, and the average wage of workers in the industry, respectively. The three average wages for each industry are estimated with CPS data. Columns 3, 5 and 7 (labeled "Emp") refer to the average employment share of skilled workers per establishment. Columns 4, 6 and 8 (labeled "Wages") refers to the average share of wages received by skilled workers per establishment. See text for more details.

Table 3 provides results analogous to those in Table 2 obtained using Nielsen's Homescan data. Here too, the share of high-skill workers employed is an increasing function of the quality of the goods produced by the firm. Consider, for example, the results we obtain using our measure of skilled workers based on the average wage of college graduates in the industry. The fraction of high-skill workers is 1.2 to 1.5 times higher in the high-quality tier when compared with the low-quality tier. The difference between the skill intensity of high- and low-quality products is lower than in the Yelp! data. This property is likely to reflect smaller differences in the quality of groceries, which are the most important category of goods in Nielsen Homescan, than in other categories such as durables.

Table 3: Quality and share of high-skill workers

Sample	#Firms	Skilled 1		Skilled 2		Skilled 3	
		Emp	Wage	Emp	Wage	Emp	Wage
Nielsen Sample	1,097	12.64	30.76	22.04	42.43	28.04	48.30
By Quality:							
Low	384	10.46	25.89	20.47	38.67	26.03	44.04
Middle	339	11.63	29.30	21.14	41.25	26.55	46.82
High	374	15.79	37.08	24.48	47.38	31.45	54.02

Notes: This table shows the average share of skilled workers per establishment for each quality tier for the food manufacturing sector. The low, middle and high quality tiers are based on the firms in the bottom third, middle third and top third of relative prices (within each product category) of firms in Nielsen Homescan data, respectively. The statistics are based on the OES survey of establishments from May 2004 to November 2007. We consider three definitions of skilled workers. In "Skill 1, 2 and 3" a worker is skilled if his/her wage exceeds the average wage of college graduates in the industry, the average wage of workers with some college education in the industry, and the average wage of workers in the industry, respectively. The three average wages for each industry are estimated with CPS data. Columns 3, 5 and 7 (labeled "Emp") refer to the average employment share of skilled workers per establishment. Columns 4, 6 and 8 (labeled "Wages") refers to the average share of wages received by skilled workers per establishment. See text for more details.

2.4 Quality and Routine Work

Acemoglu and Autor (2011) argue that the distinction between skilled and unskilled workers may be less relevant to study new forms of automation. These forms of automation, such as artificial intelligence, might replace routine tasks that are performed by high-skill workers (e.g. radiologists).

In Table 4, we study the relation between quality and the intensity of routine and non-routine work. Two patterns emerge from our data. First, the share of routine workers is lower in the production of high-quality goods when compared to that low quality goods. So quality is intensive in non-routine work. Second, the part of non-routine work that rises with quality is abstract work, not manual work.

Table 4: Quality and share of workers performing different tasks

Sample	#Firms	Routine		Non-Routine Manual		Non-Routine Abstract	
		Emp	Wage	Emp	Wage	Emp	Wage
By Quality:							
Low	384	76.66	62.78	5.24	3.36	18.10	33.87
Middle	339	80.62	62.77	2.35	1.57	17.03	35.66
High	374	69.16	51.44	7.60	3.95	23.24	44.60

Notes: This table shows the average share of workers per establishment performing each of the three following tasks: routine, non-routine manual, and non-routine abstract. We stratify the summary statistics by quality tier for establishments in the food-manufacturing sector. The low, middle and high quality tiers are based on firms in the bottom third, middle third and top third of relative prices of firms in Nielsen Homescan data, respectively. The statistics are based on the OES survey of establishments from May 2004 to November 2007. Columns 3, 5 and 7 (labeled "Emp") refer to the average employment share of routine, non-routine manual and non-routine abstract per establishment, respectively. Columns 4, 6 and 8 (labeled "Wages") refers to the average share of wages received by of routine, non-routine manual and non-routine abstract per establishment, respectively. See text for more details.

3 Quality and Income

Our second empirical fact is that the quality of the goods and services consumed by households rises with income. We document this fact using data from the CEX and the Nielsen Homescan Data. Our findings corroborate previous results in the literature. Bils and Klenow (2001) show that for a wide range of durable goods in the CEX, households with higher total expenditures consume higher-quality goods. Similarly, Faber and Fally (2018), Jaravel (2018), and Argente and Munseob (2016) show that higher-income households consume higher-quality goods. There is also a large trade literature that shows that as countries get richer, they increase the quality of what they consume (see, e.g. Verhoogen (2008) and Fieler, Eslava and Xu (2017)).

3.1 CEX

Our first data set, are durable expenditures from the CEX. Our data covers 73,000 households over the period 1980-2013. Durables are defined as categories whose items have a life that exceeds 2 years. These categories include home furnishing (e.g carpeting, curtains, mattresses, and sofas), appliances (e.g. dryers, microwaves, stoves, and radios), electronics, and vehicles. The advantage of using durables expenditures is that, given that households are unlikely to buy more than one item at a time, we can, as in Bils and Klenow (2001), use expenditures as a measure of the price paid for each item.

We estimate the quality Engel curve as follows. As in Bils and Klenow (2001), we express the unit price of an item paid by household h at time t as $x_{ht} = z_{ht}q_{ht}$, where z_{ht} is the quality-adjusted price and q_{ht} is the quality of the item. We estimate θ, the elasticity of quality with respect to income, using the following specification:

$$\ln(q_{ht}) = \beta_0 + \theta \ln(y_{ht}) + \epsilon_{ht}$$

where y denotes the income of household h in period t, and ϵ_{ht} denotes the residual.

We can rewrite this specification as

$$\ln(x_{ht}) = \beta_0 + \theta \ln(y_{ht}) + \ln(z_{ht}) + \epsilon_{ht}. \tag{1}$$

The logarithm of quality-adjusted price, $\ln(z_{ht})$, is an unobservable variable that reflects differences in prices across time and across households that are not related to the choice of quality. It can, for instance, be due to differences in shopping intensity across households which affects the prices paid for the same item. It may also be due to differences in the discounts available in different locations. As in Bils and Klenow

(2001), we include demographic controls to account for these unobservable factors that may affect prices paid. These controls include the age of the households, family size, household fixed effects, and time fixed effects.

Table 5 reports our regression results using 5 income quintile dummies, so that θ is a vector. This table shows that high-income households pay on average 80 percent more than low-income households for items of a given category. In our view, these differences are too large to be explained by price discrimination or different search intensities. For example, Aguiar and Hurst (2006) estimate that doubling the shopping frequency lowers the price paid for a given good by only 7 to 10 percent.

3.2 Nielsen Homescan

We supplement the empirical evidence on the relation between income and quality for durable goods with a second data set, the Nielsen Homescan data, that focuses on non-durable goods such as grocery products.

We compute an average price across households, $\bar{P}_{imt}$, for every item i in product module m and time t. By using this average price we ensure that differences in overall prices paid by households reflect differences in choice of item-store, rather than shopping intensity (i.e. using coupons and taking advantage of promotions). For each household, we compute the price of module m at time t as:

$$\log\left(P_{hmt}\right) = \sum_i w_{iht} \log\left(\bar{P}_{imt}\right).$$

We then estimate the following regression:

$$\log\left(P_{hmt}\right) = \beta_0 + \sum_k \beta_k 1\left(y_{ht} \in k\right) + \gamma X_{ht} + \lambda_t + \lambda_m + \varepsilon_{hmt}, \tag{2}$$

where $1\left(y_{ht} \in k\right)$ is a dummy variable equal to one if the household income is in quintile

Table 5: Prices and income: CEX

Consumer Expenditure Survey Durables	(I)	(II)
Relative to income quintile 1:		
Income quintile 2	0.205	0.197
	(0.010)	(0.010)
Income quintile 3	0.368	0.353
	(0.010)	(0.010)
Income quintile 4	0.533	0.513
	(0.010)	(0.009)
Income quintile 5 (top)	0.834	0.82
	(0.010)	(0.010)
Time fixed effects	Yes	Yes
Category fixed effects	Yes	Yes
Demographic controls	Yes	
Number of observations	824,851	824,851

Notes: This table shows the θ estimates implied by regression 1. The table reports the log-difference in price paid by each income quintile relative to the lowest income quintile. We used CEX data for the period 1980-2013. Column I includes demographic controls for age, family size and number of income earners in the household. Column II does not include demographic controls. Estimates in both columns I and II are clustered by household. See text for more details.

Table 6: Prices and income: Nielsen Homescan data

Nielsen Homescan	(I)	(II)
Relative to income quintile 1:		
Income quintile 2	0.0399	0.0398
	(0.0004)	(0.0004)
Income quintile 3	0.0911	0.0908
	(0.0004)	(0.0004)
Income quintile 4	0.151	0.150
	(0.0004)	(0.0004)
Income quintile 5 (top)	0.227	0.224
	(0.0004)	(0.0004)

Notes: This table shows the coefficients for β_k implied by regression 2. The table reports the log-difference in price paid by each income quintile relative to the lowest income quintile. We used Nielsen Homescan data for the period 2004-2010. Column I includes demographic controls for age, family size and number of income earners in the household. Column II does not include demographic controls. Estimates in both columns I and II are clustered by household. See text for more details.

k, X_{ht} denotes demographic controls (age group, employment status, size of family, and ethnicity), λ_t denotes time fixed effects, λ_m denotes product-module fixed effects, and ε_{hmt} is the error term.

Table 6 reports our estimates. The first and second columns report results without and with demographic controls, respectively. The results are similar. Households in the top quintile choose items that are 22.7 percent more expensive than households in the bottom quintile. This difference is economically large given that most of the products included in Nielsen's Homescan Data are groceries. This is an expenditure category in which price differentials are relatively small when compared with categories such as durable goods.

4 A Simple Model

In this section, we consider a simple model that is consistent with our two empirical facts. First, firms that produce higher quality goods employ a higher share of high-skill workers. In other words, quality is skill intensive. Second, the quality of the goods a household consumes rises with income. In other words, quality is a normal attribute. We then consider the implications of our model for the measurement of SBTC.

Before we delve into the details of our model, we review the key features of the canonical model used to explain the rise in the skill premium. In this model, output, Y, is produced according to the following production function:

$$Y = A\left[\alpha\,(SH)^{\rho} + (1-\alpha)L^{\rho}\right]^{1/\rho}, \tag{3}$$

where S denotes the level of SBTC, A denotes the level of HNTC, H is the supply of skilled work, and L is the supply of unskilled work.

Output is produced by firms that are competitive in product and factor markets. The optimization conditions for these firms imply that the skill premium, defined as the ratio of the wage rate of skilled, W_H, and unskilled, W_L, workers, is given by:

$$\frac{W_H}{W_L} = \frac{\alpha}{1-\alpha} S_t^{\rho} \left(\frac{H}{L}\right)^{\rho-1}. \tag{4}$$

This expression implies that changes in A have no impact on the skill premium. Computing logarithmic growth rates of the two sides of equation (4), we obtain:

$$\Delta \log \left(\frac{W_H}{W_L}\right) = \rho \Delta \log (S) + (\rho - 1)\, \Delta \log \left(\frac{H}{L}\right). \tag{5}$$

According to the estimates in Acemoglu and Autor (2010), the relative supply of effective labor of skilled workers H/L increased by about 110 percent and the skill premium increased 22 percent between 1970 and 2008. As in Acemoglu and Autor (2010), we assume that ρ equals 0.41.[8] Equation (5) implies that $\Delta \log(S) = 210.8$ percent, which corresponds to an average annual rate of SBTC of 5.5 percent.

4.1 Homogeneous households

We make some simplifying assumptions so that our model is as similar as possible to the canonical model used in the SBTC literature. These assumptions are: (i) no capital in production; (ii) a single production sector; and (iii) households consume a single unit of the consumption good, so changes in expenditures translate fully into changes in the quality of the goods consumed. These assumptions allow us to derive results without taking a stand on the form of the utility function.

Household We consider a representative household composed by the same fraction of skilled and unskilled workers present in the population. The household pools its resources and buys a single unit of a consumption good of quality q at a price $P(q)$. The household budget constraint is given by

$$P(q) = W_H H + W_L L,$$

where H and L denote high-skill and low-skill workers respectively. We treat the supply of high- and low-skill workers as exogenous and assume that workers are identical within each skill group. Household utility is given by

$$U = V(q),$$

where $V'(q) > 0$, $V''(q) \leq 0$.

[8] Estimates for ρ generally range from 0.16 (Card and Lemieux (2001) to 0.5 (Angrist (1995)).

Production Final goods are produced by competitive firms using skilled and unskilled labor according to the production function

$$Y = A\left[\alpha\left(SH\right)^{\rho} + q^{-\gamma\rho}(1-\alpha)\left(L\right)^{\rho}\right]^{\frac{1}{\rho}}, \tag{6}$$

where $\rho > 0$, $\gamma \geq 0$, and q denotes the quality of the good produced. When $\gamma = 0$ this production function is identical to the one used in the canonical model (equation (3)). Firms only produce the level or levels of quality demanded by households.

The equilibrium price of a good of quality q is given by:

$$P(q) = \frac{1}{A}\left[\alpha^{\frac{1}{1-\rho}}\left(S\right)^{\frac{\rho}{1-\rho}} W_H^{\frac{\rho}{\rho-1}} + (1-\alpha)^{\frac{1}{1-\rho}}\left(q\right)^{\frac{\gamma\rho}{\rho-1}} W_L^{\frac{\rho}{\rho-1}}\right]^{\frac{\rho-1}{\rho}}.$$

The production function (6) with $\gamma > 0$ together with the perfect-competition assumption implies two key properties. First, $P'(q) > 0$, i.e. the price of a final good is increasing in its quality. Second, quality is intensive in high-skill labor, i.e., the labor share of high-skill labor, $W_H H/\left(W_H H + W_L L\right)$ is increasing in q.

Skill Premium The firms' optimization conditions imply that the skill premium is given by:

$$\frac{W_H}{W_L} = \frac{\alpha q^{\gamma\rho}\left(S\right)^{\rho}}{(1-\alpha)}\left(\frac{H}{L}\right)^{\rho-1}. \tag{7}$$

Computing growth rates, we see that the change in the skill premium is the one obtained in the canonical model (equation (5)) plus the effect of trading up on the skill premium, which is given by the term $\gamma\rho\Delta\log(q)$ in the expression below:

$$\Delta\log\left(\frac{W_H}{W_L}\right) = \rho\Delta\log\left(S\right) + \left(\rho-1\right)\Delta\log\left(\frac{H}{L}\right) + \gamma\rho\Delta\log(q). \tag{8}$$

Equations (7-8) show that quality plays a role that is similar to that of SBTC. Other things equal, a rise in quality increases the demand for skilled labor raising the skill premium.

SBTC and HNTC Combining the household budget constraint and the firms' first-order condition we obtain the following equation:

$$A \times S = \frac{\left(\frac{W_H}{W_L}\right)^{\frac{1}{\rho}} \left(\frac{H}{L}\right)^{\frac{1-\rho}{\rho}}}{\alpha^{\frac{1}{\rho}} L \left[\frac{W_H}{W_L}\frac{H}{L} + 1\right]^{\frac{1}{\rho}}}. \tag{9}$$

Equation (9) implies that, holding constant H and L, a change in A has the same impact on the skill premium as a change in S. This equivalence holds because in our model a change in A triggers a change in q:

$$q = \left[A(1-\alpha)^{1/\rho} \left(\frac{W_H}{W_L}H + L\right)\left(\frac{W_H H}{W_L L} + 1\right)^{(1-\rho)/\rho}\right]^{1/\gamma}, \tag{10}$$

and this change in q has an effect on the skill premium that is similar to that of a rise in S (see equation (7)).

4.1.1 Quantitative results

The estimates in Acemoglu and Autor (2010) suggest that between 1970 and 2008 the skill premium increased by 25 percentage points, while the effective ratio of high-skill labor to low-skill labor increased by 110 percentage points. Below, we use our model to characterize the combinations of SBTC and HNTC that are consistent with these observed changes in W_H/W_L, H and L.

Using equation (9) to compute logarithmic growth rates, we obtain

$$\Delta A + \Delta S = \frac{1}{\rho}\left(\Delta W_H - \Delta W_L\right) + \frac{1-\rho}{\rho}\left(\Delta H - \Delta L\right) - \Delta L + \frac{1}{\rho}\Delta\left(1 + \frac{W_H H}{W_L L}\right). \tag{11}$$

Following Acemoglu and Autor (2010), we set $\rho = 0.41$. Given this value of ρ, we can compute the right-hand side of equation (11), which is equal to 1.92 percent on an annual basis. The left-hand side of the equation gives us the combinations of ΔA and ΔS that match the right-hand side.

Table 7 reports results for both our model and the canonical model for different combinations of HNTC and SBTC. We choose γ equal to 1.15 so that for our benchmark calibration (row 3) our model is consistent with the Bils-Klenow estimates for Δq (3.8 percent per year).

Table 7: Trading up and Canonical Model

ΔA	ΔS	Cumulative $\Delta(W_H/W_L)$ (percent)	
		Trading-up model	Canonical model
0.00	0.00	-46	-65
0.87	0.00	-25	-65
0.87	1.05	25	-48

Notes: All numbers in percentages. ΔA and ΔS are reported on an annual basis. $\Delta(W_H/W_L)$ is the value for the entire sample.

Line 1 corresponds to the case where there is no HNTC or SBTC. In the canonical model, the skill premium falls by 65 percent. In our model, this fall is only 46 percent. The reason for this result is that in our sample more workers became skilled over time, creating a rise in income that induces households to trade up, expanding the demand for high-skill labor and increasing the skill premium.

In line 3, we set the rate of SBTC to zero and choose the annual rate of HNTC to match Fernald's (2014) estimates for TFP growth for the period 1970 to 2008 (0.87 percent per year). Our model predicts a 25 percent fall in the skill premium because the rise in the demand for skilled workers is not strong enough to overcome the increase in the relative supply of skilled workers. In this calibration, HNTC accounts for 30 percent of the rise in the skill premium while in the canonical model it accounts for zero percent of this rise.[9]

[9]In the absence of HNTC and SBTC, the skill premium falls by 46 percent (row 1 of column 4). HNTC and SBTC together raise the change in the skill premium from -46 to +25 percent. With

Row 3 reports results for our benchmark specification. We set the annual rate of HNTC to match Fernald's (2014) estimates and choose the annual growth rate of SBTC so that our model is consistent with the observed rise in the skill premium. The required annual rate of SBTC is 1.05 percent, a much more plausible number that the 5.5 percent required by the canonical model to match the observed rise in the skill premium. In the canonical model, the configuration of HNTC and SBTC considered in row 3 results in a 48 percent fall in the skill premium. Comparing rows 2 and 3 we see that even though in our benchmark calibration the rate of SBTC is only 1.05 percent per year, the presence of SBTC is essential to produce a rise in the skill premium.

The future of the skill premium Our model can shed some light on the future evolution of the skill premium. We start by forecasting the fraction of high-skill workers in 2026 using data from the U.S. Department of Education. To simplify, we treat this fraction as exogenous. High-skill workers are defined as those with educational attainment equal to some college or associate degree, bachelor's and advanced degree. A linear regression fits well the data from 1992 to 2016, so we base our projections on this regression. We forecast a rise in the fraction of high-skill workers from 62 percent in 2008 to 71 percent in 2026.

To isolate the impact of this increase in the supply of skilled workers, we first abstract from technical progress. We calibrate our model with the same value of γ used in the benchmark case (1.15). Our results are reported in the first row of Table 8. The canonical model implies a moderation in income inequality with the skill premium falling by 21 percent. In contrast, our model predicts a decline of only 15 percent in the skill premium. This smaller decline in the skill premium occurs because the preponderance of high-skill workers in the labor force means higher household income which boosts the demand for quality, increasing the demand for skilled labor.

HNTC and no SBTC, the skill premium falls by 25 percent. So the fraction of the rise in the skill premium accounted for by HNTC is $[-25-(-46)]/(25-(-46)) = 30$ percent.

Table 8: Forecasting the Skill Premium: Trading up and Canonical Model

ΔA	ΔS	Cumulative $\Delta(W_H/W_L)$ (percent)	
		Trading-up model	Canonical model
0.0	0.0	-14	-21
0.8	0.0	25	-21

Next, we use the forecast of the rate of change in HNTC produced by Fernald (2016), which is 0.8 percent per year. Our results are reported in row 2 of Table 8. Our model implies that the skill premium will increase by 25 percent between 2008 and 2026, the same that occurred between 1970 to 2008. These results reflect the fact that, according to our projections, the fraction of high-skill workers will increase by less in the future than in the 1970-2008 period.

4.2 Heterogenous Households

In the model described above we make several simplifying assumptions to derive analytical results about the impact of trading up on the skill premium. In particular, there is only one quality level produced in equilibrium. So the model cannot, by construction, match the cross-sectional relation between quality and the share of high-skill workers used in production.

In this subsection, we extend the model by considering an economy where there is more than one consumption good produced in equilibrium. This version of our model has two types of households. The first type has only low-skill workers and the second type only high-skill workers. These households receive different income levels and, as a result, they consume goods of different quality.

Production Household type $j \in \{L, H\}$ consumes goods of quality q_j which are produced according to

$$Y_{q_j} = A\left[\alpha\left(SH_j\right)^{\rho} + q_j^{-\gamma\rho}\left(1-\alpha\right)L_j^{\rho}\right]^{1/\rho}.$$

The price of a good of quality q_j is

$$P_{q_j} = \frac{1}{A}\left[\alpha^{1/(1-\rho)}S^{\rho/(1-\rho)}W_H^{\rho/(1-\rho)} + q_j^{\frac{\gamma\rho}{\rho-1}}(1-\alpha)^{1/(1-\rho)}W_L^{\rho/(1-\rho)}\right]^{(\rho-1)/\rho}.$$

We assume perfect labor mobility which implies that the skill premium is identical in both the high- and low-quality production sector. The first-order conditions for competitive output producers imply that the skill premium in each sector j is given by:

$$\frac{W_H}{W_L} = \frac{\alpha}{1-\alpha}\left(q_j^{\gamma} \times S\right)^{\rho}\left(\frac{H_j}{L_j}\right)^{\rho-1}.$$

For future reference, we note that using the expression for the skill premium implies that total expenditure in each of the goods produced is given by

$$P_{q,j}Y_{q,j} = W_L L_{q,j}\left[1 + \frac{W_H}{W_L}\left(\frac{H_{q,j}}{L_{q,j}}\right)\right].$$

Households There is a measure $\Gamma(H)$ of skilled workers and $\Gamma(L) = 1 - \Gamma(H)$ of unskilled workers. The supply of skilled work measured in efficiency units is: $H = H_{q,L} + H_{q,H}$. Similarly, the supply of unskilled workers in efficiency units is: $L = L_{q,L} + L_{q,H}$.

The maximization problem of high-skill households is:

$$\max U = V(q_H),$$

subject to

$$P(q_H) = W_H H.$$

Similarly, the maximization problem of low-skill households is:

$$\max U = V(q_L),$$

subject to

$$P(q_L) = W_L L.$$

Equilibrium Using the budget constraints of low- and high–skill workers, we obtain:

$$L_{q,L}\left[1+\frac{W_H}{W_L}\left(\frac{H_{q,L}}{L_{q,L}}\right)\right] = L,$$

$$L_{q,H}\left[1+\frac{W_H}{W_L}\left(\frac{H_{q,H}}{L_{q,H}}\right)\right] = \frac{W_H}{W_L}H.$$

As in the previous section we look for combinations of SBTC and HNTC that are consistent with the observed changes in the skill premium and aggregate changes in L and H.

To solve this two-sector model we proceed as follows. We guess $H_{q,L}/L_{q,L}$. The budget constraint of low-skill workers implies a value for $L_{q,L}$. Using this value, we can compute $H_{q,L}$ and $H_{q,H}$. We can then compute the ratio of the two qualities produced:

$$\frac{q_H^{\gamma}\left(\frac{W_H}{W_L}\frac{H_H}{L_H}+1\right)^{\frac{\rho-1}{\rho}}\Gamma_H}{q_L^{\gamma}\left(\frac{W_H}{W_L}\frac{H_L}{L_L}+1\right)^{\frac{\rho-1}{\rho}}\Gamma_L} = \frac{W_H}{W_L}\frac{H}{L}.$$

Given the observed skill premium, the overall measures of high- and low-skill workers, and the equilibrium sectoral allocation, these identify the ratio of qualities, q_H/q_L. With this ratio we then verify that the skill premium is identical in both sectors.

Before commenting on the implications for the measurement of the SBTC, we note that we can compute the share of high-skill workers in labor income $(W_H H/\left(W_L H + W_H L\right))$ in the production of the high- and low-quality good. Importantly, the model has no free parameters that allows us to target these ratios, so it is of interest to see whether the model comes close to the observed numbers in the data. The ratio of these shares is 2.4 in the beginning of the sample and 2.5 in the end of the sample. We also compare the share of high-skill labor ($H/(H+L)$) used in the production of the high- and the low-quality good. The ratio of these shares is 2.8 in the beginning of the sample and 3.6 in the end of the sample. These values somewhat higher than those reported in Table 2 for the ratio of highest quality ($$$) and the lowest quality ($).

Once we find the equilibrium sectoral allocation, we use the fact that each consumer purchases one unit of the good and thus, $Y_{q,j} = \Gamma(j)$ from which it follows that for each sector the following equation holds

$$\Delta A - \gamma \Delta q_H = \frac{\rho - 1}{\rho} \Delta \left(1 + \frac{W_H}{W_L} \frac{H_H}{L_H} \right) + \Delta \Gamma_H - \Delta \left(\frac{W_H}{W_L} H \right).$$

The right hand side is a given number (given the data and the equilibrium sectoral allocation). Hence given different values of ΔA we can recover Δq for each of the two sectors. Then from the skill premium equation,

$$\frac{W_H}{W_L} = \frac{\alpha}{1-\alpha} \left(q_j^{\gamma} \times S \right)^{\rho} \left(\frac{H_j}{L_j} \right)^{\rho - 1}.$$

we can recover the consistent value of ΔS.

We re-do the analysis for the same configurations of γ, ΔA and ΔS used in Table 7. The results are almost identical, showing that our results are robust to introducing the form of income heterogeneity embodied into this model. The reason for this robustness is that in our calibration the price-quality schedule is close to linear, so that producing two levels of quality instead of one does not affect substantively the properties of the equilibrium. The linearity of the price-quality schedule comes from two sources. First, the value of $\rho\,(0.41)$ is close to 0.5 which makes the exponent $(\rho-1)/\rho$ in equation (4.2) close to one. Second, the value of γ (1.15) is also close to one.

An interesting question is: how do real wages of low- and high-skill workers evolve over time in our benchmark calibration? We deflate wages with the price level evaluated at a constant quality level, the one consumed in the beginning of the sample. These quality levels are denoted by q_{0H} and q_{0L} for high-skill and low-skill households, respectively. The real wage for high- and low-skill workers are $W_H/P(q_{0H})$ and $W_L/P(q_{0L})$, respectively. The prices $P(q_{0H})$ and $P(q_{0L})$ change between the beginning and end of the sample as a result of endogenous changes in wages and the two types of technical progress, SBTC and HNTC. The annual rate of change in prices are as follows: $\Delta P(q_{0L}) = -1.3$ percent, and $\Delta P(q_{0H}) = -4.7$ percent. Since W_L is the numeraire, the rate of change in the real wage of the low-skill household is the symmetric of the price change: $\Delta W_L/P(q_{0L}) = -1.3$ percent. The rate of change in the real wage of the high-skill household is: $\Delta W_H/P(q_{0H}) = 5.3$ percent. Since the price of quality falls more for the high-skill workers than for the low-skill workers, the ratio of the real wages of high- and low-skill grows by more than the skill premium. In other words, the divergence in real wages across skill groups exceeds the rise in the skill premium.

5 Quantity and Quality

In our benchmark model, whenever a household receives additional income it devotes it fully to increasing the quality of the good it consumes. In this section, we explore how the impact of trading up changes when the household buys both one unit of a quality good, available in a continuum of different qualities, and a basic good (C) with a fixed level of quality.

We assume that the utility function takes the form used by Bils and Klenow (2001):

$$U = \frac{C^{1-1/\sigma}}{1-1/\sigma} + \nu \frac{q^{1-1/\sigma_q}}{1-1/\sigma_q}. \tag{12}$$

The budget constraint of the representative household is

$$P(q) \times 1 + C = HW_H + LW_L.$$

The basic good is produced according to the following production function

$$Y_c = A \left[\alpha \left(SH_c\right)^\rho + (1-\alpha) \left(L_c\right)^\rho\right]^{\frac{1}{\rho}}.$$

The quality good is produced according to

$$Y_q = A \left[\alpha \left(SH_q\right)^\rho + q^{-\gamma\rho}(1-\alpha) \left(L_q\right)^\rho\right]^{\frac{1}{\rho}}.$$

The labor market clearing conditions are

$$H = H_c + H_q,$$
$$L = L_c + L_q.$$

The goods market clearing conditions are

$$Y_q = 1,$$
$$Y_c = C.$$

We calculate the annual rate of change in SBTC that rationalizes the rise in the skill premium over the period 1970-2008. We use the same annual rate of HNTC (0.83) and the same value of ρ (0.41) used in the benchmark model. The value of σ is set to one and the ratio σ_q/σ is chosen to match the estimates of Bils and Klenow (2001), which are 0.76.

We search for combinations of α, ν, and γ that allow us to match three targets. The first is the annual growth rate for q_t estimated by Bils and Klenow (3.8 percent). The second is the average share of quality goods in expenditure (41 percent). The third is the ratio of the share of skilled labor in income in high quality goods versus low quality goods from Table 3:

$$\frac{\frac{W_H H_q}{W_H H_q + W_L L_q}}{\frac{W_H H_c}{W_H H_c + W_L L_c}} = 1.6$$

Here we are using the basic good as the low quality good and the quality good as the high-quality good.

We find that the level of SBTC necessary to rationalize the skill premium is 1.4 percent. This value is in between the value obtained for the benchmark model (1.05 percent) and the canonical model (5.5 percent), but closer to the value implied by the benchmark model.

The intuition for these results is as follows. A rise in income produces less trading up because some of the additional income is spent on the basic good. As a result, the effect a rise in income on the skill premium is weaker than in the benchmark model. There is, however, another effect. The utility function (12) implies that the quality good is a superior good, so its share in spending rises as income rises, expanding the demand for skilled workers and contributing to increasing the skill premium. This effect is absent in our benchmark model. In that model the share of income spent in the high quality good cannot rise because it is already 100 percent.

6 Conclusions

In this paper, we show empirically that as income rises, households trade up to higher quality goods and that these goods are intensive in skilled labor. As a result, the demand for high-skill labor rises, increasing the skill premium. This trading-up phenomenon amplifies the effect of skill-biased technical change by creating an endogenous rise in the demand for skilled workers.

In the canonical model, technical change has to be skill biased to produce a rise in the skill premium. In our model, any factor that raises income increases the skill premium. This idea has important implications for the future evolution of the skill premium. Growth in income that is not accompanied by an increase in the supply of skilled workers is likely to increase the skill premium even in the absence of skill-biased technological change.

Acemoglu and Autor (2011) argue that the distinction between skilled and unskilled workers may be less relevant to study forms of automation such as artificial intelligence which may replace routine tasks performed by high-skill workers. Motivated by this argument, we study the relation between quality and the intensity of routine and non-routine work. We find that low-quality goods are intensive in routine work and that high-quality goods are intensive in non-routine, abstract work. So, as households trade up, we expect to see a rise in the relative wages of non-routine, abstract workers.

The choice of the quality of goods and services consumed by households has implications not just for the skill premium but also for changes in the overall wage distribution. We find that quality increases lead to a relative decline in the demand for routine workers. According to the estimates of Autor and Dorn (2013), workers in the 10th to 50th percentile of the 1980 wage distribution have jobs with a high share of routine tasks. So the wages of these workers are likely to decline relative to the wages of workers on the bottom and top of the wage distribution.

Taken as a whole, our evidence suggests that considering endogenous changes in the quality of consumption is an important avenue for future work on the dynamics of income inequality.

7 Appendix: Capital Deepening and the Skill Premium

Krusell, Ohanian, Ríos-Rull, and Violante (2000), which we refer to as KORV, argue that capital deepening associated with investment-specific technical progress explains much of the rise in the skill premium. To re-examine their analysis, we adopt their functional form for the production function. Output with quality q is produced according to the following nested CES function

$$Y_q = K_S^{\gamma}\left[\alpha\left(\lambda K_E^{\sigma} + (1-\lambda)H^{\sigma}\right)^{\frac{\rho}{\sigma}} + q^{-\gamma\rho}(1-\alpha)\left(L\right)^{\rho}\right]^{\frac{1-\gamma}{\rho}}, \tag{13}$$

where K_S is the stock of structures and K_E is the stock of equipment.

To retain the one-sector character of the model, we assume that investing I_{qt} units of output with quality q_t yields $P(q_t)$ units of installed capital. The capital accumulation equation takes the form

$$\begin{aligned} K_{St+1} &= P_{qt}I_{qSt} + (1-\delta)K_{St}, \\ K_{Et+1} &= P_{qt}I_{qEt} + (1-\delta)K_{Et}, \end{aligned}$$

where I_{qSt} and I_{qEt} denote the investment in structures and equipment, respectively.

The market clearing condition for output is:

$$Y_{qt} = 1 + I_{qSt} + I_{qEt}/z_t,$$

where $1/z_t$ is the relative price of equipment. Recall that the representative household buys one unit of a good with quality q_t.

We use the wage of low-skill workers as the numeraire,

$$W_L = 1,$$

but for clarity we retain the symbol W_L in the derivations below.

Output producers are competitive in goods and factors markets. Profit maximization implies the following first-order conditions

$$P_q \gamma K_S^{\gamma-1} \left[\alpha \left(\lambda K_E^{\sigma} + (1-\lambda) H^{\sigma}\right)^{\frac{\rho}{\sigma}} + q^{-\gamma\rho}(1-\alpha)\left(L\right)^{\rho}\right]^{\frac{1-\gamma}{\rho}} = R_S,$$

$$P_q(1-\gamma) K_S^{\gamma} \left[\alpha \left(\lambda K_E^{\sigma} + (1-\lambda) H^{\sigma}\right)^{\frac{\rho}{\sigma}} + q^{-\gamma\rho}(1-\alpha)\left(L\right)^{\rho}\right]^{\frac{1-\gamma-\rho}{\rho}} \times \tag{14}$$
$$\alpha \left(\lambda K_E^{\sigma} + (1-\lambda) H^{\sigma}\right)^{\frac{\rho-\sigma}{\sigma}} \frac{\lambda K_E^{\sigma}}{K_E} = R_E,$$

$$P_q(1-\gamma) K_S^{\gamma} \left[\alpha \left(\lambda K_E^{\sigma} + (1-\lambda) H^{\sigma}\right)^{\frac{\rho}{\sigma}} + q^{-\gamma\rho}(1-\alpha)\left(L\right)^{\rho}\right]^{\frac{1-\gamma-\rho}{\rho}} \times \tag{15}$$
$$\alpha(1-\lambda) \left(\lambda K_E^{\sigma} + (1-\lambda) H^{\sigma}\right)^{\frac{\rho-\sigma}{\sigma}} \frac{H^{\sigma}}{H} = W_H,$$

$$P_q(1-\gamma) K_S^{\gamma} \left[\alpha \left(\lambda K_E^{\sigma} + (1-\lambda) H^{\sigma}\right)^{\frac{\rho}{\sigma}} + q^{-\gamma\rho}(1-\alpha)\left(L\right)^{\rho}\right]^{\frac{1-\gamma-\rho}{\rho}} \times \tag{16}$$
$$q^{-\gamma\rho}(1-\alpha)\left(L\right)^{\rho-1} = W_L.$$

where R_S and R_E are the rental rates on structures and equipment, respectively.

These first-order conditions imply that the price of a good with quality q is

$$P_q = \frac{R_S^{\gamma} \left[\alpha^{\frac{1}{1-\rho}} \left(\lambda^{\frac{1}{1-\sigma}} R_E^{\frac{\sigma}{\sigma-1}} + (1-\lambda)^{\frac{1}{1-\sigma}} W_H^{\frac{\sigma}{\sigma-1}}\right)^{\frac{\sigma-1}{\sigma}\frac{\rho}{\rho-1}} + (1-\alpha)^{\frac{1}{1-\rho}} q^{\left(\frac{\gamma\rho}{\rho-1}\right)} \left(W_L\right)^{\frac{\rho}{\rho-1}}\right]^{\left(\frac{\rho-1}{\rho}\right)(1-\gamma)}}{\gamma^{\gamma}(1-\gamma)^{1-\gamma}}. \tag{17}$$

Equations (15) and (16) imply that the skill premium is given by

$$\frac{\alpha(1-\lambda)\left(\lambda K_E^{\sigma} + (1-\lambda) H^{\sigma}\right)^{\frac{\rho-\sigma}{\sigma}} H^{\sigma-1}}{(1-\alpha) q^{-\gamma\rho} L^{\rho-1}} = \frac{W_H}{W_L}. \tag{18}$$

Combining equations (14) and (15), we obtain

$$\frac{(1-\lambda)}{\lambda}\left(\frac{K_E}{H}\right)^{-\sigma} R_E K_E = W_H H. \tag{19}$$

Combining equations (18) and (19), the skill premium can be written as

$$\frac{\alpha(1-\lambda)^{\frac{\rho}{\sigma}}}{(1-\alpha)} \times \frac{\left[1+R_E K_E/\left(W_H H\right)\right]^{\frac{\rho-\sigma}{\sigma}}}{q^{-\gamma\rho}}\left(\frac{H}{L}\right)^{\rho-1} = \frac{W_H}{W_L}. \tag{20}$$

If we abstract from the impact of quality by assuming that q is constant and equal to one, we obtain the same expression for the skill premium equation used in KORV.

7.1 Some Analytics

Recall that σ is the parameter that governs the elasticity of substitution between capital and high-skill workers. To see the effects of the presence quality on the point estimates forσ, it is useful to log-linearize equation (20),

$$(\rho-\sigma)\frac{1}{1+W_H H/\left(R_E K_E\right)}\left(\widehat{K}_E-\widehat{H}\right)+(\rho-1)\left(\widehat{H}-\widehat{L}\right)+\gamma\rho\widehat{q}=\widehat{W_H}-\widehat{W_L}. \tag{21}$$

where $\hat{x}$ denotes the logarithmic growth rate of x. Solving equation (21) for σ, we obtain

$$\begin{aligned} \sigma \quad = \quad & \frac{1+W_H H/\left(R_E K_E\right)}{\widehat{K}_E-\widehat{H}} \times & (22) \\ & \left[\frac{\rho(\widehat{K}_E-\widehat{H})}{1+W_H H/\left(R_E K_E\right)}+(\rho-1)(\widehat{H}-\widehat{L})-\widehat{W_H}+\widehat{W_L}+\gamma\rho\widehat{q}\right]. & (23) \end{aligned}$$

Capital-skill complementarity requires $\sigma \leq \rho$. How are these estimates affected by trading up? The answer to this question depends on the value of $\gamma\rho\widehat{q}[1+W_H H/\left(R_E K_E\right)]/(\widehat{K}_E-\widehat{H})$. In the KORV data $\widehat{K}_E-\widehat{H}>0$. Since $\rho>0$, an increase in $\widehat{q}$ implies that the right-hand side is overall a higher number. Since $\sigma<0$, then the degree of capital skill complementarity required to match the same empirical facts is reduced.

Trading up also affects the point estimates of ρ, the parameter that governs the degree of substitutability between unskilled workers and the composite good of equipment capital and skilled workers. To see this effect, note that the value of ρ, as a function of a given value of σ can be expressed as

$$\rho = \frac{\frac{\sigma\left(\widehat{K}_E-\widehat{H}\right)/\left(1+\frac{W_H H}{R_E K_E}\right)+\left(\widehat{W_H}+\widehat{H}-\widehat{W_L}-\widehat{L}\right)}{\left(\widehat{K}_E-\widehat{H}\right)/\left(1+\frac{W_H H}{R_E K_E}\right)+\left(\widehat{H}-\widehat{L}\right)}}{1+\frac{\widehat{\gamma q}}{\left(\widehat{K}_E-\widehat{H}\right)/\left(1+\frac{W_H H}{R_E K_E}\right)+\left(\widehat{H}-\widehat{L}\right)}}. \tag{24}$$

The change in the level of quality affects only the denominator reducing the value of ρ. As a result, the degree of substitutability of unskilled labor and the composite good of equipment and skilled worker $(1/(1-\rho))$ falls.

In this analysis we were holding constant the value of σ when analyzing the effects of quality on the measurement of ρ and vice versa when analyzing the effect of quality for the measurement of σ. Naturally, both of these estimates can change as a results of quality. We thus proceed by jointly estimating these two parameters.

7.2 Estimation

In this section, we estimate the production function using the approach proposed by Polgreen and Silos (2008).[10] In this approach, the posterior distribution is obtained by combining a prior distribution for the vector of parameters with a measurement-error-based likelihood function for the data.

The estimation is based on three conditions. The first is the equation for the labor share in income

[10]We are extremely grateful to Pedro Silos for kindly sharing with us his code and for various consultations.

$$\frac{W_L L}{P_q Y_q}+\frac{W_H H}{P_q Y_q} = (1-\gamma) Y_q^{-\rho/(1-\gamma)} \left[\alpha(1-\lambda)\left(\lambda K_E^{\sigma}+(1-\lambda) H^{\sigma}\right)^{\frac{\rho-\sigma}{\sigma}} H^{\sigma}+q^{-\gamma\rho}(1-\alpha) L^{\rho}\right].$$

The second condition is the equation for the ratio of labor income of skilled and unskilled agents,

$$\frac{W_H H}{W_L L}=\frac{\alpha(1-\lambda)\left(\lambda K_E^{\sigma}+(1-\lambda) H_q^{\sigma}\right)^{\frac{\rho-\sigma}{\sigma}} H^{\sigma}}{q^{-\gamma\rho}(1-\alpha) L_q^{\rho}}.$$

The third condition equates the rate of return to investing in structures and equipment.

$$\begin{aligned}
& \gamma K_S^{\gamma-1}\left[\alpha\left(\lambda K_E^{\sigma}+(1-\lambda) H^{\sigma}\right)^{\frac{\rho}{\sigma}}+q^{-\gamma\rho}(1-\alpha)(L)^{\rho}\right]^{\frac{1-\gamma}{\rho}}+(1-\delta) \\
= \quad & (1-\gamma) K_S^{\gamma}\left[\alpha\left(\lambda K_E^{\sigma}+(1-\lambda) H^{\sigma}\right)^{\frac{\rho}{\sigma}}+q^{-\gamma\rho}(1-\alpha)(L)^{\rho}\right]^{\frac{1-\gamma-\rho}{\rho}} \times \\
& \alpha\left(\lambda K_E^{\sigma}+(1-\lambda) H^{\sigma}\right)^{\frac{\rho-\sigma}{\sigma}} \lambda K_E^{\sigma-1}+(1-\delta) z_{t-1}/z_t.
\end{aligned}$$

Estimating ρ and σ We begin by replicating the analysis of Polgreen and Silos (2008) and estimate ρ and σ in a model without quality choice. The resulting estimates are $\rho = 0.4470, \sigma = -0.3871$.

We now estimate ρ and σ in a model with quality choice. Since we do not have a time series for Δq, we consider a constant trend in quality at the annual growth rate estimated by Bils and Klenow (2001). We obtain the following estimates: $\rho = 0.2485$ and $\sigma = -0.3730$. As suggested by our analytical results, incorporating quality choice into the model implies a fall in the degree of substitutability of unskilled labor and the composite good of equipment and skilled worker $(1/(1-\rho))$. This elasticity of substitution between unskilled labor and capital falls from $1/(1-0.4470) = 1.81$ to $1/(1-0.2485) = 1.33$.