Review of Economic Dynamics 17 (2014) 842–859

Contents lists available at ScienceDirect

Review of Economic Dynamics

www.elsevier.com/locate/red

Gender gaps across countries and skills: Demand, supply and the industry structure ☆

Claudia Olivetti [a], Barbara Petrongolo [b,c,*]

[a] Boston University, United States
[b] Queen Mary University, United Kingdom
[c] CEP (LSE), United Kingdom

ARTICLE INFO

Article history:
Received 20 August 2012
Received in revised form 7 March 2014
Available online 1 April 2014

JEL classification:
E24
J16
J31

Keywords:
Gender gaps
Skills
Demand and supply
Industry structure

ABSTRACT

The comovement between gender gaps in hours and wages across countries and skills reveals the presence of net demand forces shaping gender differences in labor market outcomes. This paper links the rich pattern of variation in gender gaps to the process of structural transformation. Based on a stylized, multi-sector equilibrium model, we illustrate that the gender bias in labor demand can be decomposed into measurable within- and between-industry components. Using comparable micro data across countries, we find that international differences in the industry structure explain more than eighty percent of the overall variation in labor demand between the U.S. and all other countries in our sample, and roughly one third of the overall cross-country variation in wage and hours gaps.

1. Introduction

Gender gaps vary widely across countries, and across levels of human capital within countries. To give an example, in the U.S., the U.K., and other countries in northern Europe, the gender wage gap is either rising with levels of education, or roughly flat, while in southern Europe gender wage penalties are largest among the unskilled. Large variations in wage gaps are accompanied by substantial variation in the corresponding gaps in hours per head.[1] In particular, gender gaps in hours everywhere fall with levels of education, but such gradient is highest in southern Europe and Ireland, where employment rates of unskilled women are lowest. This pattern of variation yields a positive cross-country correlation between the unskilled-to-skilled difference in the wage gap and the corresponding difference in the gap in hours per head (see Fig. 1). Based on a canonical labor supply and demand framework, positive co-variation in price and quantity differentials reveals the presence of net demand forces shaping gender differences in labor market outcomes across skills and countries.

☆ We wish to thank the editor, Gianluca Violante, and an anonymous referee for their insightful suggestions. We also thank Claudia Goldin and Steve Pischke for their comments on earlier versions of this paper, Craig Riddel for invaluable help with the Canadian LFS, and the NSF and ESRC for financial support (NSF Grant SES-0820127, Olivetti; ESRC Grant RES-000-22-4114, Petrongolo). Briana Ballis, Maria Burtseva and Michael Harris provided excellent research assistance. Helpful comments from seminar participants at Bocconi University, Boston University, the 2009 Summer Meetings of the Econometric Society, the 2009 ESPE Conference, the 2010 Eale/Sole Conference, the 2010 Research on Money and Markets Conference and the ESOP-ISF Conference on Gender and Households, and the 2011 NBER Winter Meetings (Labor Studies) are also gratefully acknowledged.
* Corresponding author.
[1] Throughout the paper, gender gaps are defined as ratios (or log differences) of male to female outcomes.

http://dx.doi.org/10.1016/j.red.2014.03.001

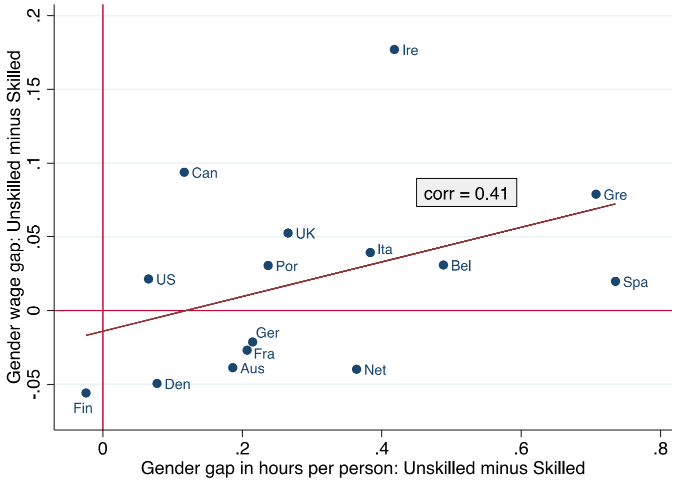

Notes. Samples, sources and wage and hours gaps are defined in notes to Table 1.

Fig. 1. Gender gaps in wages and hours per person: unskilled-to-skilled differences.

Table 1
Gender gaps in wages and hours-to-population ratios.

Countries	(1)	(2)	(3)	(4)
	(log) Gender gap in wages		(log) Gender gaps in hours per head	
	Unskilled	Skilled	Unskilled	Skilled
U.S.	0.286	0.265	0.321	0.255
Canada	0.243	0.149	0.315	0.198
U.K.	0.288	0.236	0.602	0.337
Finland	0.181	0.237	0.196	0.221
Denmark	0.103	0.153	0.269	0.191
Germany	0.227	0.249	0.600	0.386
Netherlands	0.211	0.251	0.850	0.486
Belgium	0.153	0.122	0.801	0.313
Austria	0.234	0.273	0.559	0.373
Ireland	0.259	0.082	0.825	0.407
France	0.174	0.201	0.495	0.288
Italy	0.082	0.043	0.758	0.375
Spain	0.202	0.182	1.090	0.355
Portugal	0.208	0.178	0.418	0.181
Greece	0.204	0.125	1.127	0.419

Notes. The skilled are those with a college degree; the unskilled are all others. Columns 1 and 2: Values displayed are gender differences in log wages by country and skill, using population weights. Columns 3 and 4: Values displayed are gender differences in log(hours/population) by country and skill, using population weights. Sample: men and women aged 25–54, excluding military, students, and self-employed. Sample period: 1994–2001, except for Canada (1997–2004), Finland (1996–2001) and Austria (1995–2001). Source: CPS (King et al., 2010), Canadian LFS, and ECHPS.

In this paper we exploit the skill dimension of gender gaps to draw new insights on the factors shaping gender inequalities across countries. We link the rich pattern of variation in gender gaps to the process of structural transformation, and investigate the role of differences in the industry composition of employment in shaping the structure of labor demand. Insofar as different industries employ a different mix of labor inputs, defined by gender and skill, we expect the industry structure to have an impact on gender gaps across countries and skills. In particular, our analysis shows that differences in the service share are an important determinant of the cross-country variation in women's labor market outcomes.

We document gender biases in labor demand using comparable micro data across countries, and find that these tend to be larger, and to display more pronounced cross-country variation, for the unskilled than for the skilled. Based on a stylized, multi-sector equilibrium model, we show that the gender bias in labor demand can be decomposed into within- and between-industry components. Within-industry forces reflect differences in gender and skill intensities within sectors, including skill-biased technical change, changes in prices of other inputs, outsourcing, or discrimination. Between-industry forces reflect differences in the sectoral composition of the economy, where different sectors may have different skill and gender intensities, and may be driven by differences in product demand, in sector-specific productivity growth, in the extent of marketization of home production, or international trade.

We find that in most countries the gender gap in labor demand, relative to the U.S., is higher for the unskilled than for the skilled, with the important qualification that in southern Europe labor demand differences with respect to the U.S.

are much larger than elsewhere, and these are driven by especially large gender gaps in wages *and* working hours for the unskilled in southern Europe.

According to our model-based decomposition, the between-industry component explains more than eighty percent of the overall variation in labor demand between the U.S. and all other countries in our sample. The industry structure thus plays an important role overall in shaping international differences in gender outcomes. However, there is noticeable variation in its importance. In Denmark, Ireland, France and Italy the between-industry component explains the whole variation in labor demand relative to the U.S. In most other countries, it explains a large fraction of the total, ranging from one third in Spain to three quarters in Austria. The within-industry component is also sizable – in both absolute and proportional terms – in Spain, Portugal and Greece, consistent with larger gender biases in labor demand for the unskilled in these countries, both between and within industries. In a further decomposition, we relate the within-industry component to international differences in the occupational structure within industries. Although not very large, the size of the between-occupation component is positively correlated to the size of the between-industry component, implying that countries that have an industry structure that favors a certain gender/skill mix also tend to have an occupational structure that is relatively more favorable to the same mix.

We finally assess the between-industry component of gender gaps by showing the counterfactual correlation between the unskilled-to-skilled wage gap and the corresponding hours gap that would be observed in equilibrium, having corrected both wage and hours gaps for the between-industry component of labor demand. For plausible values of labor supply elasticities, this adjustment explains around one third of the comovement between wage and hours gaps. Moreover we show that the international variation in the service share is the key between-industry force that drives our findings, and we relate such variation to country-level indicators of institutions, attitudes towards female work, and uneven productivity growth.

The relationship between structural transformation and women's involvement in the labor market has been noted as far back as Fuchs (1968). The idea is that production of goods and services is relatively intensive in the use of "brawn" and "brain", respectively, and as men and women may have different endowments of these factors, the historical growth in the service sector would disproportionately attract women into the labor market. A similar point is made by Goldin (1995, 2006), who notes that the secular expansion of services has made available jobs that were physically less demanding and more respectable for women joining the labor force than typical jobs in factories, and Olivetti (2014), who documents the relationship between women's role in the labor market and the process of structural transformation since the late 19th century in a number of countries. As the decline in manufacturing and the parallel rise in services may be staggered across countries, these ideas have consequences for the international variation in female labor market outcomes. Our paper makes a contribution in this direction.

By looking at the role of the industry structure in shaping gender gaps across skills and countries, this paper brings together two strands of literature. First, there is a literature studying the causes of the variation in the gender gap across countries. Work by Blau and Kahn (1996, 2003) emphasizes the role of international differences in overall wage dispersion: if women tend to have on average lower wage characteristics than men, higher overall inequality would map these differences into a wider gender pay gap. Our previous work (Olivetti and Petrongolo, 2008) stresses the process of selection into paid work and concludes that if working women tend to have relatively high-wage characteristics, low female employment rates become consistent with low wage gaps simply because low-wage women are less likely to feature in the observed wage distribution. We contribute to this literature by uncovering the skill dimension of gender inequalities, and relating the variation in gender gaps across skills and countries to the industry structure.

Second, this paper builds on a large literature on the impact of structural transformation and, more generally, technological progress, on the wage and employment structure. In particular, our paper is related to work by Ngai and Pissarides (2008), emphasizing the role of uneven productivity growth across sectors in the secular reallocation of hours of work in the U.S. from manufacturing to services, and Autor and Dorn (2013), showing how the reduced usage of routine tasks following the adoption of new technologies may have driven employment and wage polarization in the U.S.[2] In both studies, the rise in the share of service industries (or service occupations) stems from the interaction between technological progress in manufacturing and poor substitutability between manufacturing and service output in consumption. Moreover, in Ngai and Pissarides (2008) technological improvements in market production, relative to home production, further boost the (market) service share via marketization of household tasks.

Technological progress and the rise in services may in turn have consequences for the gender structure of labor demand. Heathcote et al. (2010) relate the bulk of the rise in female hours in the U.S. since the early 1980s to a gender-biased demand shift, and Ngai and Petrongolo (2014) highlight the impact of the rise of services on trends in female hours and relative wages. In an international perspective, Rogerson (2007, 2008) and Ngai and Pissarides (2011) relate differences in hours between continental Europe and the U.S. to the smaller weight of the service sector in Europe, and Rendall (2013) and Akbulut (2011) emphasize different implications for men's and women's hours. These papers highlight the marketization of services that have close substitutes in home production as a key force driving the rise in services and variation in market hours. Our approach complements these studies along two main dimensions. First, we introduce both gender and skill

[2] See Goos et al. (forthcoming) for evidence on other countries, and Black and Spitz-Oener (2010) for a study of the impact of routinization on the gender wage gap in West Germany.

dimensions in the analysis of the labor market effects of structural transformation. If skilled and unskilled women tend to be over-represented in different industries, special attention should be paid to the impact of the industry structure on female labor market outcomes across the skill distribution. Second, we emphasize international differences in both gender gaps and the industry structure, and illustrate the role of the between-industry component of labor demand in shaping the international variation in gender outcomes.[3]

The paper is organized as follows. Section 2 presents some key facts on the variation of gender gaps across skills and countries. Section 3 proposes a multi-sector model to decompose the variation in labor demand into measurable between- and within-industry components. Section 4 presents the results of the model-based decomposition and highlights the role of the service share in driving our main findings. Section 5 concludes.

2. Facts on gender gaps by skill level

This section presents some key facts on gender gaps by levels of education for the U.S., Canada and thirteen European countries. These are: U.K., Finland, Denmark, Germany, Netherlands, Belgium, Austria, Ireland, France, Italy, Spain, Portugal and Greece. For the U.S. we use data from the March Current Population Survey (CPS) for 1995–2002, where each year's survey contains detailed information on the previous year's labor market variables. This choice of sample period is made to ensure consistency with European data, extracted from the European Community Household Panel Survey (ECHPS), which is only available from 1994 to 2001, and provides contemporaneous information on labor market variables. For Canada we use data from the March Labor Force Survey (LFS) for 1997–2004.[4]

While the data may differ in structure – for the U.S. and Canada we use repeated cross-sections, for Europe we have an unbalanced panel – the information we exploit from these data is consistent across countries. We select individuals aged 25–54 who are not in full-time education, retired, military, or self employed. Weekly hours for the U.S. are usual weekly hours worked in the previous year. For Canada and Europe, we use information on usual hours in the survey week. Hourly wages are obtained for the U.S. as gross wage and salary income in the previous year, divided by annual hours. For Canada, we use directly available information on current gross hourly earnings. For Europe, wages are obtained by dividing current gross monthly wage and salary earnings by actual hours worked, as a measure of usual earnings would not be available in the ECHPS. Our core sample includes individuals with positive earnings and hours. As the definition as well as the adoption of part-time work varies widely across countries, we do not restrict the analysis to full-time workers.

Information on educational attainment is only available in the ECHPS by broad categories, i.e. less than upper secondary high school, upper secondary school completed, and higher education. These correspond to ISCED 0–2, 3, and 5–7, respectively. We thus attempt to reproduce this threefold distinction for the U.S. and Canada, where available categories of education are 15 and 7, respectively. For the U.S. and Canada the low education group includes all individuals who have not completed 12th grade, the middle group includes those who have completed 12th grade but do not have a college degree, and the high-education group includes those who have completed a college degree. Education shares in the population for each country are reported in Table B1 in Appendix B.

Our analysis is based on a twofold skilled/unskilled distinction, thus we need to reorganize the three educational categories available into two groups. An obvious solution is to merge the mid-education group to either the low- or high-education group. This is equivalent to treating secondary school graduates as either pure dropout equivalents or pure college equivalents. To determine which one of these options is more appropriate, as in Katz and Murphy (1992), we regress mean wages for high school graduates by year, country and gender on mean wages for dropouts and college graduates, plus controls for year, country and gender. The regression results show that a person with a high school degree is equivalent to a total of 0.983 of a high-school dropout (with a standard error of 0.058), and −0.014 of a person with a college degree (with a standard error of 0.028). We thus merge the low- and middle-education groups to form our unskilled labor group, and the skilled group only includes college graduates. This classification also has the advantage to define as skilled a group whose qualifications are measured relatively consistently across countries. Table B2 in Appendix B shows wage bill shares by gender for the skilled and unskilled, as defined above.

Table 1 shows gender gaps in wages and hours per person by skill. Columns 1 and 2 display wage gaps. In Nordic countries and a group of continental European countries including Germany, Netherlands, Austria and France, the gender wage gap is higher for the skilled than for the unskilled, though the proportional difference is stronger in Nordic countries than elsewhere. In the rest of the sample the wage gap is instead higher for the unskilled. While in the U.S. such difference is rather small, at least in proportional terms, it becomes quite sizable in other countries, and especially in Canada, Ireland, Italy and Greece. Columns 3 and 4 show gender differences in the (log of the) hours to population ratio. In all countries except Finland the gender gap in hours per head falls substantially with the level of education, but the gradient is much

[3] While we focus on the effect of the industry structure on the demographic composition of employment, it should be recognized that changes in female labor supply may in turn have an impact on the industry mix. However, existing evidence suggests this should not be a major issue. In particular, Lee and Wolpin (2006) conclude that the growth in the service industry resulted almost entirely from demand-side factors associated with technical change, and that supply-side factors are neutral with respect to relative sector growth. Other than this, it should be noted that if causality were running from female labor supply to the industry structure, the resulting correlation between quantities and wages would be the opposite of what we observe.

[4] Information on wages and earnings is first included in the Canadian LFS in 1997, and in order to use eight survey years for Canada as for most other countries we extend the corresponding sample until 2004.

stronger in Belgium, Ireland and Southern Europe than elsewhere. Interestingly, countries differ widely in their unskilled hours gap, with much lower variation in the skilled hours gap.

This rich variation in gender gaps can be broadly summarized by looking at the correlation between the skill differential in wage and hours gaps. Fig. 1 plots the difference between the unskilled and the skilled wage gap (i.e. the difference between columns 1 and 2 in Table 1) against the difference between the unskilled and the skilled gaps in hours per head (i.e. the difference between columns 3 and 4). The correlation between them is positive, equal to 0.41. There is clearly no outlier that drives this correlation, and excluding each country in turn from the sample gives correlation estimates ranging from 0.32 (excluding Finland) to 0.52 (excluding Canada). Positive comovements of quantity and price differentials clearly reveal the presence of *net* demand factors shaping the variation in gender gaps across skills and countries.

Before exploring the nature of cross-country differences in labor demand, it should be noted that the demographic groups considered are characterized by very different employment rates, and one may worry about the way in which different patterns of employment selection across genders, skills and countries may affect our results. Imagine, for the sake of the argument, that in a country with low female participation it would be socially acceptable for an educated woman to take a skilled job, but it would not seem proper for an uneducated woman to take an unskilled job as a cleaner or waitress, unless she is "forced" by economic conditions in her household. As a result, fewer uneducated women would work, and those who do would be negatively selected on household characteristics, and may have low-wage characteristics themselves, resulting in higher wage gaps at the bottom of the wage distribution. This hypothetical outcome, although observationally equivalent to some of the patterns observed in Table 1, would not be driven by differences in demand forces, but simply by differences in the quality composition of the employed workforce in different countries.

Below we use a very simple method for controlling for selection, which consists in imputing wages to the non-employed based on their observable characteristics, and then estimating median wage gaps on the resulting enlarged wage distribution. By relying on median, as opposed to mean, wage gaps, the only information that is exploited about imputed wages is their position with respect to the median of the potential wage distribution, not the actual imputed level.[5] Our imputation follows two alternative rules. With the first rule, we impute wages below the median (by gender and skill) to all those who are unemployed, as opposed to nonparticipants, and we leave the potential wages of nonparticipants as missing. The underlying idea is that the unemployed are receiving wage offers (if any) below their reservation wage, while the employed have received at least one wage offer above their reservation wage. At given reservation wages, the unemployed have lower potential wages than the observed wages of the employed, and are thus assigned an imputed wage below the median. With the second rule, we assign wages below the median to non working individuals whose partners have total income in the bottom quartile of their gender/skill-specific distribution, based on the assumption of positive assortative mating along wage attributes. The correlation between wage and hours gaps stays firmly positive once we control for selection into paid work using either the first imputation rule (0.50) or the second imputation rule (0.47). Hence we find no evidence at all that employment selection behavior could explain the observed variation of gender gaps by skill.

3. A multisector model

3.1. The economy

We propose an equilibrium model of a multi-sector economy, in which the demand for each labor input is driven by both the intensity of its use in each industry and the industry structure, as different industries may employ various inputs in varying proportions. The model economy outlined below corresponds to a country in our data set. We assume that each country is a closed-economy, ruling out cross-country flows of inputs and outputs.

We consider an economy with J industries, and assume that output in each industry j, Q_j, is produced by a combination of skilled and unskilled labor, denoted by S_j and U_j respectively, according to the following CES production function:

$$Q_j = \left[\alpha_j S_j^{\frac{\phi-1}{\phi}} + (1-\alpha_j) U_j^{\frac{\phi-1}{\phi}}\right]^{\frac{\phi}{\phi-1}}, \tag{1}$$

where α_j is a technology parameter representing the relative weight of skilled labor in industry j, and ϕ denotes the elasticity of substitution between skilled and unskilled labor.[6]

We assume that skilled and unskilled labor are each described by CES aggregators of female and male labor inputs:

$$S_j = \left[\beta_{Sj}(B_{MSj} M_{Sj})^{\frac{\sigma-1}{\sigma}} + (1-\beta_{Sj})(B_{FSj} F_{Sj})^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{2}$$

$$U_j = \left[\beta_{Uj}(B_{MUj} M_{Uj})^{\frac{\sigma-1}{\sigma}} + (1-\beta_{Uj})(B_{FUj} F_{Uj})^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{3}$$

where M_{Sj}, M_{Uj}, F_{Sj} and F_{Uj} represent the four types of labor, B_{MSj}, B_{FSj}, B_{MUj} and B_{FUj} are the associated labor-augmenting technological parameters, and β_{Sj} and β_{Uj} denote the share of work activities performed by men from each skill group in

[5] See Olivetti and Petrongolo (2008) for a formal discussion of this methodology.

[6] See Bound and Johnson (1992) and Katz and Autor (1999).

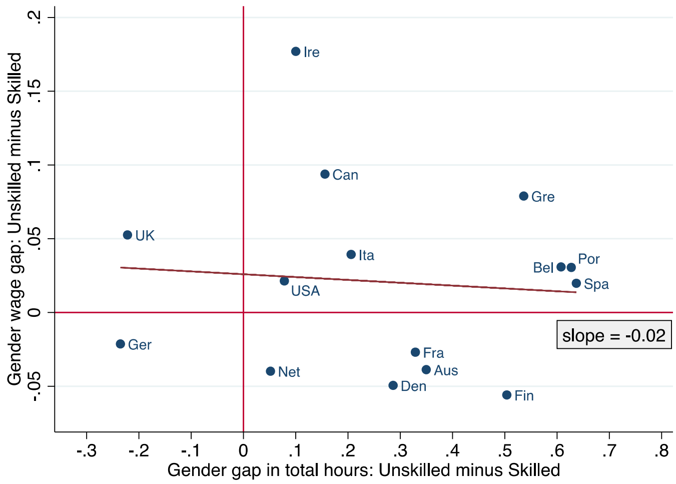

Notes. Wage gaps are defined in notes to Table 1. Hours gaps are gender differences in log hours by country and skill, using population weights.

Fig. 2. Gender gaps in wages and total hours: unskilled-to-skilled differences.

each industry. The parameter σ represents the elasticity of substitution between male and female labor, assumed to be constant across skills and industries. This assumption will be relaxed below.

Aggregate output is given by $Q = \sum_j A_j Q_j$, where A_j denotes total factor productivity in industry j. We express demand for industry-j output relative to the reference industry-r output by the unit price-elasticity function

$$\frac{Q_j}{Q_r} = \theta_j P_j^{-1}, \tag{4}$$

where P_j denotes the price of Q_j relative to the numeraire good Q_r, and θ_j is a demand shifter capturing shocks to relative product demand for industry-j output.

3.2. Cross-country differences in relative labor demand

Our first goal is to quantify the gender and skill dimensions of differences in labor demand across countries. To this purpose, we relate gender wage gaps to (measurable) gaps in labor demand and labor supply. The wage gap for skill group i ($i = S, U$) is denoted by $\Delta w_i \equiv \ln(\frac{W_{Mi}}{W_{Fi}})$, where W_{Mi} and W_{Fi} are gender specific wages. The gap in labor supply is denoted by $\Delta h_i \equiv \ln(\frac{M_i}{F_i})$, with $M_i = \sum_j M_{ij}$ and $F_i = \sum_j F_{ij}$, and the corresponding gap in labor demand is denoted by $\tilde{\beta}_i$. Under perfect competition in the labor market, all inputs are paid their marginal productivity and wages are equalized across sectors. This assumption delivers the following expression for Δw_i:

$$\Delta w_i = \tilde{\beta}_i - \frac{1}{\sigma}\Delta h_i, \tag{5}$$

stating that the wage gap for a given skill is simply given by the gap in labor demand, net of the corresponding gap in labor supply. Gaps in labor demand for the skilled and unskilled are in turn given by

$$\tilde{\beta}_S = \frac{1}{\sigma}\ln\left(\frac{\sum_j \theta_j^\sigma \tilde{\alpha}_j^\sigma \beta_{Sj}^\sigma B_{MSj}^{\sigma-1} S_j^{1-\sigma}}{\sum_j \theta_j^\sigma \tilde{\alpha}_j^\sigma (1-\beta_{Sj})^\sigma B_{FSj}^{\sigma-1} S_j^{1-\sigma}}\right) \tag{6}$$

and

$$\tilde{\beta}_U = \frac{1}{\sigma}\ln\left(\frac{\sum_j \theta_j^\sigma (1-\tilde{\alpha}_j)^\sigma \beta_{Uj}^\sigma B_{MUj}^{\sigma-1} U_j^{1-\sigma}}{\sum_j \theta_j^\sigma (1-\tilde{\alpha}_j)^\sigma (1-\beta_{Uj})^\sigma B_{FUj}^{\sigma-1} U_j^{1-\sigma}}\right), \tag{7}$$

respectively (see Appendix A.1 for derivation). Note that $\tilde{\beta}_S$ and $\tilde{\beta}_U$ are expressed as functions of between-industry components, represented by the distribution of consumer demand across industries (θ_j), and within-industry components, represented by different gender/skill intensities in each industry. The latter include all industry-specific technology parameters, B_{MSj}, B_{FSj}, B_{MUj}, B_{FUj}, β_{Sj} and β_{Uj}; skill-specific inputs in each industry, S_j and U_j; and skill-specific wage bill shares, denoted by $\tilde{\alpha}_j$ and $1 - \tilde{\alpha}_j$ respectively.

Using expression (5), the cross-skill difference in the gender wage gap is:

$$\Delta w_U - \Delta w_S = (\tilde{\beta}_U - \tilde{\beta}_S) - \frac{1}{\sigma}(\Delta h_U - \Delta h_S), \tag{8}$$

and is driven by skill differences in gender gaps in demand ($\tilde{\beta}_U - \tilde{\beta}_S$) and supply ($\Delta h_U - \Delta h_S$).

Consider next how the double difference in (8) varies across countries. To ease this comparison, assume for simplicity that all countries share a common σ, and that the only factors that vary across countries are relative demands and supplies of labor inputs. Thus the (triple) difference in wages across genders, skills and countries can be expressed as

$$\Delta_C(\Delta w_U - \Delta w_S) = \Delta_C(\tilde{\beta}_U - \tilde{\beta}_S) - \frac{1}{\sigma}\Delta_C(\Delta h_U - \Delta h_S), \tag{9}$$

where Δ_C indicates the differential between a generic country C in our sample and the U.S.

A way to simply assess the importance of demand differences across countries consists in working out the comovement of wage differentials and hours differentials in a scenario in which the structure of labor demand is equalized across countries, i.e. $\Delta_C(\tilde{\beta}_U - \tilde{\beta}_S) = 0$. In this case equation (9) implies a negative cross-country relationship between $(\Delta w_U - \Delta w_S)$ and $(\Delta h_U - \Delta h_S)$, with a slope equal to $-1/\sigma$. In fact, we plot $(\Delta w_U - \Delta w_S)$ against $(\Delta h_U - \Delta h_S)$ in Fig. 2, and note that the underlying relationship is essentially flat, with a slope of -0.02. Under reasonable values for the elasticity of substitution σ around 2.5,[7] such slope should be -0.4 in the absence of relative demand differences. Thus we note that the cross-country variation in wage and hours gaps can only be rationalized by underlying net relative demand differences.[8]

3.3. Equilibrium gender gaps

To obtain equilibrium gaps in wages and hours we close the model summarized in Eq. (8) by introducing a labor supply relationship. We consider the simplest labor supply model, in which hours of work respond to wages with a common elasticity $\eta > 0$:

$$\Delta w_U - \Delta w_S = (\delta_U - \delta_S) + \frac{1}{\eta}(\Delta h_U - \Delta h_S), \tag{10}$$

where the δ's are skill-specific labor supply shifters. Combining (8) and (10) gives the following equilibrium conditions for wage and hours differences, respectively:

$$\Delta w_U - \Delta w_S = \frac{\sigma}{\sigma + \eta}(\tilde{\beta}_U - \tilde{\beta}_S) - \frac{\eta}{\sigma + \eta}(\delta_U - \delta_S); \tag{11}$$

$$\Delta h_U - \Delta h_S = \frac{\sigma\eta}{\sigma + \eta}\big[(\tilde{\beta}_U - \tilde{\beta}_S) - (\delta_U - \delta_S)\big]. \tag{12}$$

Eqs. (11) and (12) will be used in Section 4.2 to quantify the impact of measurable variation in labor demand on the equilibrium skill differential in wage gaps, for realistic values of σ and η.

3.4. Between- and within-industry components of labor demand

Our next goal is to decompose cross-country differences in labor demand into between- and within-industry components. Using (5), the cross-country difference in the gender gap in demand for skill group $i = S, U$ is:

$$\Delta_C\tilde{\beta}_i = \Delta_C\Delta w_i + \frac{1}{\sigma}\Delta_C\Delta h_i. \tag{13}$$

Note that, conditional on the wage gap $\Delta_C\Delta w_i$, a higher elasticity of substitution σ translates a given hours gap $\Delta_C\Delta h_i$ into a smaller labor demand gap $\Delta_C\tilde{\beta}_i$. The intuition is that when male and female labor inputs are highly substitutable, a small change in relative demand can generate large fluctuations in relative hours. Thus, for a given gap in hours, the larger σ, the smaller the underlying gap in labor demand differentials.

The between-industry component of (13) can be obtained by differentiating (6) and (7) with respect to θ_j, for the skilled and the unskilled, respectively. This gives:

$$\Delta_C\tilde{\beta}_i^{between} = \left[\frac{\sum_j M_{ij,C}W_{Mi,C}\Delta_C\theta_j/\theta_{j,C}}{M_{ij,C}W_{Mi,C}} - \frac{\sum_j F_{ij,C}W_{Fi,C}\Delta_C\theta_j/\theta_{j,C}}{F_{ij,C}W_{Fi,C}}\right], \tag{14}$$

for $i = S, U$.

All terms on the right hand side of Eq. (14) are measurable. In particular, θ_j is given by industry j's share of national revenue, which in turn equals industry j's share of the total wage bill, and $M_{ij}W_{Mi}$ and $F_{ij}W_{Fi}$ are wage bills by gender and skill.

7 Appendix A will provide evidence on this.

8 Note that Figs. 1 and 2 plot different hours variables on the x-axis. Preliminary evidence given in Fig. 1 is based on gaps in hours per head, in order to factor in cross-country variation in educational attainment. Fig. 2 uses instead gaps in total hours, consistent with predictions from the model of this section. The interpretation of the different correlations obtained is that cross-country variation in the educational attainment of the population raises further the correlation between wage and hours gaps from about zero (Fig. 2) to 0.4 (Fig. 1).

Using notation y for wage bill shares, expression (14) can be rewritten as

$$\Delta_C \tilde{\beta}_i^{between} = \frac{\sum_j \bar{y}_{Mij} \Delta_C y_j}{\bar{y}_{Mi}} - \frac{\sum_j \bar{y}_{Fij} \Delta_C y_j}{\bar{y}_{Fi}}, \tag{15}$$

where y_j denotes industry j's share of the total wage bill, y_{Mij} (y_{Fij}) denotes the wage bill share of males (females) of skill i in industry j, y_{Mi} (y_{Fi}) denotes the wage bill share of males (females) of skill i in the economy, and upper bars denote averages between country C and the U.S. (see Appendix A.1 for derivation).

Expression (15) shows that the between-industry component of labor demand differences is independent of the elasticity of substitution σ. However, the total difference in (13) depends on σ, and specifically it falls with σ whenever $\Delta_C \Delta h_i > 0$. Thus, the relative weight of the between-industry component increases with σ for $\Delta_C \Delta h_i > 0$ and the opposite holds for $\Delta_C \Delta h_i < 0$.

The corresponding within-industry component can be simply obtained as:

$$\Delta_C \tilde{\beta}_i^{within} = \Delta_C \tilde{\beta}_i - \Delta_C \tilde{\beta}_i^{between}. \tag{16}$$

While the industry structure is the main focus of our discussion, the above decomposition can be used to look into other dimensions of the composition of labor demand. In particular, we consider role of the occupation structure in shaping labor demand patterns across countries, and this links to a growing task-based approach to changes in labor demand (see Acemoglu and Autor, 2011, for an extensive survey). As changes in the occupation structure may take place within industries, we further decompose the within-industry component in (16) into a between-occupation and a within-occupation component. The between-occupation component hinges on differences in the occupation structure within each given industry, and is given by:

$$\Delta_C \tilde{\beta}_i^{occ} = \sum_j \bar{y}_j \left[\frac{\sum_q \bar{y}_{Mijq} \Delta_C y_{jq}}{\bar{y}_{Mij}} - \frac{\sum_q \bar{y}_{Fijq} \Delta_C y_{jq}}{\bar{y}_{Fij}} \right], \tag{17}$$

where y_{jq} denotes the wage bill share of occupation q in industry j, y_{Mijq} (y_{Fijq}) denotes the wage bill share of males (females) of skill i and occupation q in industry j, y_{Mij} (y_{Fij}) denotes the wage bill share of males (females) of skill i in industry j, and again upper bars denote averages between country C and the U.S.

3.5. Extensions

The above framework makes the rather extreme assumption that male and female inputs are equally substitutable across skills and across industries. Here we relax this assumption allowing, first, such elasticity to vary across skills, and, second, to vary across industries.

Assume that male and female inputs have skill-specific elasticity of substitution σ_i, which is constant across industries. In this case gender differences in labor demand for the skilled and the unskilled are derived as in (6) and (7), respectively, having replaced the common elasticity parameter σ with the skill-specific one, σ_S or σ_U, respectively. As the between-industry component of labor demand differences is independent of the elasticity parameter, this is still represented by expression (15). Allowing for skill-specific elasticity thus simply rescales the total gender bias in labor demand for each skill group, while leaving its between-industry component unchanged.

Allowing for further variation in the male/female elasticity of substitution across industries has more substantial consequences because it precludes a closed-form solution for the gender wage gap. To see this, it can be shown that the gender wage gap for the skilled would be implicitly defined by the following expression:

$$\Delta w_S = \frac{1}{\bar{\sigma}_S} \ln\left(\frac{\sum_j \theta_j^{\sigma_{Sj}} \tilde{\alpha}_j^{\sigma_{Sj}} \beta_{S_j}^{\sigma_{Sj}} B_{MSj}^{\sigma_{Sj}-1} S_j^{1-\sigma_{Sj}} W_{MS}^{(\bar{\sigma}_S - \sigma_{Sj})}}{\sum_j \theta_j^{\sigma_{Sj}} \tilde{\alpha}_j^{\sigma_{Sj}} (1-\beta_{Sj})^{\sigma_{Sj}} B_{FSj}^{\sigma_{Sj}-1} S_j^{1-\sigma_{Sj}} W_{FS}^{(\bar{\sigma}_S - \sigma_{Sj})}} \right) - \frac{1}{\bar{\sigma}_S} \Delta h_S, \tag{18}$$

where σ_{Sj} denotes the elasticity parameter for the skilled in industry j, and $\bar{\sigma}_S$ denotes its average across industries (see Appendix A.2 for derivation). The main difference between Eqs. (18) and (5) is represented by two extra factors in wages on the right-hand side of (18), $W_{MS}^{(\bar{\sigma}_S - \sigma_{Sj})}$ and $W_{FS}^{(\bar{\sigma}_S - \sigma_{Sj})}$, which serve as weights for industry-specific labor demand terms. In particular, sectors in which male and female labor are less easily substitutable bear a higher weight in the wage gap, as a labor demand shock in a low-elasticity sector would impact relative wages more than relative hours. Note that the expression for $\tilde{\beta}_S$ obtained from (6) can still provide a reasonable approximation to the gender difference in labor demand when the variation in the sector-specific elasticity of substitution is small enough. A similar expression can be obtained for Δw_U.

The corresponding formula for the between-industry component of cross-country labor demand differences for group $i = S, U$ is given by:

$$\Delta_C \tilde{\beta}_i^{between} = \frac{\sum_j \frac{\sigma_{ij}}{\bar{\sigma}_i} \bar{y}_{Mij} \Delta_C y_j}{\bar{y}_{Mi}} - \frac{\sum_j \frac{\sigma_{ij}}{\bar{\sigma}_i} \bar{y}_{Fij,C} \Delta_C y_j}{\bar{y}_{Fi}}, \tag{19}$$

Table 2
Decomposition of labor demand differences.

Countries	(1) Triple differences in labor demand $\sigma = 2.5$ (benchmark)	(2) Between-industry component	(3) Triple differences in labor demand $\sigma = 1.5$	(4) Triple differences in labor demand $\sigma = 3.5$	(5) Between-occupation component (within industry)
Canada	12.88	6.32	14.77	12.06	1.22
U.K.	−8.12	−3.46	−16.30	−4.61	0.25
Finland	15.63	7.43	28.58	10.08	−0.95
Denmark	3.91	4.41	8.94	1.75	−0.54
Germany	−16.98	4.27	−24.98	−13.55	0.34
Netherlands	−7.76	−0.05	−8.65	−7.38	0.58
Belgium	17.07	6.80	28.95	11.98	1.11
Austria	6.57	5.22	15.24	2.85	0.79
Ireland	14.24	14.79	15.41	13.74	1.11
France	2.89	2.70	10.09	−0.20	−0.69
Italy	5.88	10.26	8.74	4.65	1.07
Spain	20.37	7.19	34.65	14.25	0.82
Portugal	28.93	18.61	47.87	20.81	1.00
Greece	23.22	12.98	34.73	18.28	1.09
% between industry	82.1		49.2	115.0	

Notes. Column 1 reports triple differences in labor demand (obtained by evaluating equation (13) for $\sigma = 2.5$) and column 2 reports their between-industry component (obtained from Eq. (15)). The percentage explained by the between-industry component (bottom of column 1) is obtained as the ratio between the cross-country sum of terms in column 2 and the cross-country sum of terms in column 1. Columns 3 and 4 report triple differences in labor demand for $\sigma = 1.5$ and $\sigma = 3.5$, respectively, with the corresponding percentages explained by the between industry component reported at the bottom. Column 5 reports within-industry, between-occupation components (from Eq. (15)). See notes to Table 1 for sample and source.

where each industry-specific term is weighted by the ratio of the industry-specific elasticity relative to the average (see Appendix A.2 for derivation).

4. Results

4.1. Model-based decomposition of labor demand differences

The previous section has shown that the between- and within-industry components of international differences in labor demand can be easily evaluated using data on wage bill shares by gender and skill and available estimates of the elasticity of substitution between male and female labor. We next perform this decomposition based on a ten-fold industry classification for each country. Industries are: primary and utilities[9]; manufacturing; construction; transport, storage, post and telecommunications; wholesale and retail trade, hotels and restaurants; financial intermediation, insurance and real estate; education; health and social work; other services; public administration and defense.

The results are reported in Table 2, for three alternative values of $\sigma = \{1.5, 2.5, 3.5\}$. Hamermesh (1993) reviews various studies on the elasticity of substitution between male and female labor and suggests values of σ of 2 for the U.K. and 2.3 for Australia. For the U.S., Weinberg (2000) provides an estimate of σ of 2.4, which is remarkably similar to the values obtained for Australia and the U.K., Johnson and Keane (2013) obtain estimates ranging from 1.85 to 2.2, and Acemoglu et al. (2004) obtain a slightly higher estimate around 3. On our dataset, we estimate σ to be about 2.2, by regressing (log) wage gaps on (log) hours gap, instrumented by their lag, and controlling for year, country, and skill effects, as well as country-specific skill effects. The case $\sigma = 2.5$ roughly coincides with the mean of existing estimates, and we consider this as our benchmark. In the sensitivity analysis we also consider the cases $\sigma = 1.5$ and $\sigma = 3.5$, which represent lower and upper bounds respectively among available estimates.

Column 1 in Table 2 reports triple differences in labor demand between each country in our sample and the U.S. for $\sigma = 2.5$. These are obtained as $\Delta_C\tilde{\beta}_U - \Delta_C\tilde{\beta}_S$, where $\Delta_C\tilde{\beta}_i$ is given by expression (13). Figures reported show that, in all countries except the U.K., Germany and the Netherlands, the gender gap in labor demand, relative to the U.S., is higher for the unskilled than for the skilled. Moreover, this gap is relatively higher in southern Europe. Column 2 reports its between industry component, as given by expression (15). This is everywhere positive, except in the U.K. and the Netherlands, where the triple difference is also negative. Overall, the between industry component explains 82% of the total cross-country variation in labor demand. This is obtained as the ratio between the cross-country sum in column 2 and the cross-country sum in column 1, and reported at the bottom of column 1.[10]

[9] These include: agriculture, hunting, forestry, fishing, mining, quarrying, electricity, gas and water supply.

[10] 82% of the overall variation in labor demand across countries is driven by differences in the industry structure and, specifically, 35% relates to differences among the skilled and 47% relates to differences among the unskilled. Gender gaps for both the skilled and the unskilled are thus affected (in the same direction) by between-industry differences in demand. However such differences seem to be stronger for the unskilled than for the skilled.

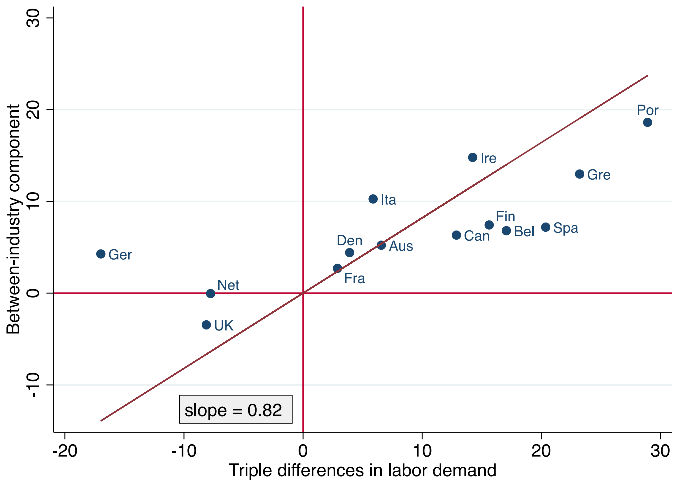

Notes. Triple differences in labor demand correspond to values reported in column 1 of Table 2. Their between-industry components correspond to values reported in column 2. See notes to Table 2 for details.

Fig. 3. The between industry component of labor demand differences.

These decomposition results can be best grasped visually in Fig. 3, which plots the between-industry component (from column 2 in Table 2) against the total difference to be explained (from column 1). The straight line has slope 0.82, representing the proportion of the total difference that is explained by the between-industry component in the whole sample. In countries to the right of the vertical axis the gender gap in labor demand, relative to the U.S., is higher for the unskilled than for the skilled, while the opposite is true to its left. In countries above the 0.82 line, the weight of the between industry component is higher than the average, and the opposite happens below it. The plot highlights noticeable cross-country differences in the importance of the between-industry component. In particular, in Denmark, Ireland, France and Italy the between-industry component explains (close to) the whole variation in labor demand, or even more than that, implying a negative within-industry component. In most other countries the between-industry component explains an important fraction of the total, ranging from 35% in Spain to 79% in Austria. Finally, in the Netherlands and Germany the between-industry component is either nil or goes in the opposite direction of the total difference.

Columns 3 and 4 in Table 2 report triple differences in labor demand based on alternative values of the elasticity of substitution between male and female labor (σ). The measured differences fall with σ in all countries except the U.K. and Germany. This is because, as implied by Eq. (13), higher σ assigns a lower weight to triple differences in hours, which are positive in all countries, except the U.K. and Germany. However, alternative values of σ leave the between-industry component (15) unaffected, thus this is still given by values reported in column 2. As a consequence, the proportion of the total variation in labor demand that is explained by the between-industry component rises with σ, from about one half for $\sigma = 1.5$ to more than the total for $\sigma = 3.5$.[11]

While differences in the industry structure explain overall a substantial portion of the international variation in labor demand, there are also important within-industry differences. In particular, the within-industry component is dominant in Germany, the Netherlands and – to a lesser extent – the U.K., in which the total difference in labor demand to be explained is negative. The interpretation is that the gender gap in labor demand is relatively lower for the unskilled in these three countries than in the U.S., and this happens in most industries. The within-industry component is also sizable in southern Europe and, most-notably, Spain, revealing a generalized (within-industry) gender bias in labor demand for the unskilled in these countries. This is consistent with previous evidence on gender gaps in unemployment rates by Azmat et al. (2006). Their findings suggest that institutions that compress wages or mostly affect groups with weaker labor market attachment (like firing costs and fixed-term contracts) may magnify gender differences in labor demand in southern Europe, and these institutions tends to be more binding for the labor market outcomes of the unskilled.

Column 5 further explores within-industry differences in labor demand and reports differences which, *within* each industry, can be explained by the occupation structure, according to expression (17). For the sake of cell size, we consider three broad occupation groups, and namely: managers, professionals and technical occupations; middle-skill occupations, including clerical and sales occupations, skilled manual and laborer occupations; and service occupations, including all jobs that involve helping, caring for, or assisting others. This is the three-fold classification of occupations emphasized by Acemoglu and Autor (2011) in order to illustrate polarization of labor demand. Figures in column 5 show that the between-occupation component is (proportionally) higher in Canada, Belgium, Austria, Ireland and southern Europe, and its

[11] The impact of σ on the total variation in labor demand can be understood going back to Eq. (13), implying that $\Delta_C \tilde{\beta}_i$ becomes more positive for lower values of σ whenever $\Delta_C \Delta h_i > 0$. This condition is satisfied for all countries (except the U.K.) for the unskilled, and for only about half of the countries for the skilled. Thus $\sum_C \Delta_C \tilde{\beta}_U$ falls with σ, while the change in $\sum_C \Delta_C \tilde{\beta}_S$ is ambiguous (and in practice $\sum_C \Delta_C \tilde{\beta}_S$ does not change much).

Table 3
Robustness tests: Varying elasticity of substitution between male and female labor Percentages of total demand difference explained by the between-industry component.

Varying elasticity by skill:	
$\sigma_U = 1.5$; $\sigma_S = 3.5$	41.8
$\sigma_U = 2.5$; $\sigma_S = 5$	71.0
Varying elasticity by industry group:	
$\sigma_{manuf} = 1.5$; $\sigma_{serv} = 3$	103.8
Varying elasticity by skill and industry group:	
$\sigma_{U,manuf} = 1.5$; $\sigma_{U,serv} = 3$	80.5
$\sigma_S = 2.5$	

Figures reported are obtained as ratios between the cross-country sum of between-industry components and the cross-country sum of total demand differences. See notes to Table 1 for sample and source.

importance tends to be positively correlated to that of the between-industry component. In other words, countries where the industry structure favors a certain labor input, relative to the U.S., tend to have an occupational structure that favors the same input.

The quantitative exercise above rests on the assumption of a constant elasticity of substitution between male and female labor across skills and sectors. We next relax this assumption by letting σ vary by skill and/or industry. While this is an interesting exercise, we are effectively constrained in our robustness tests by the limited available evidence on σ at the skill or industry level.

There is only scant evidence on how substitutability of male and female labor varies by skill. To our knowledge, the only available evidence is provided by Acemoglu et al. (2004), who find an elasticity of substitution in the range 2.5 to 4 between all women and men with a high-school degree in the U.S., and an elasticity ranging from 4 to 10 between all women and male college graduates. Here we take on board the evidence that σ would be higher among the skilled than the unskilled, with the important qualification that the estimates of Acemoglu et al. (2004) refer to the 1940–1960 period, when gender substitutability in production may have been quite different from the 1990s for several reasons, and do not distinguish between skilled and unskilled women. Based on this evidence, we consider two cases. In the first case we choose ($\sigma_U = 1.5$; $\sigma_S = 3.5$), which is centered around our benchmark value. In the second case we set ($\sigma_U = 2.5$; $\sigma_S = 5$), which mimics overall higher values of σ found by Acemoglu et al. (2004). For each case, we report in Table 3 the proportion of the total cross-country variation that is explained by the between-skill component (obtained as the cross-country sum of between-industry components and the cross-country sum of total demand differences). When $\sigma_U = 1.5$ and $\sigma_S = 3.5$, the between-industry component explains about 42% of the total (row 1), and this proportion rises to 71% when $\sigma_U = 2.5$ and $\sigma_S = 5$ (row 2). As argued above, given the international pattern of unskilled hours gaps, it is mostly variation in σ_U that affects our measure of total demand differences and their between-industry components. Specifically, low substitutability between unskilled female and unskilled male labor implies a larger demand differential and a smaller role for the between-industry component across countries. However, for reasonable values of σ_U, the between-industry component still explains a sizable portion of labor demand differences.

We finally let σ vary by industry. For simplicity we classify industries into "services" (including all service industries and the public sector), and "goods" (including the primary sector, manufacturing and construction). While there is no direct evidence on the variation of the elasticity of substitution between male and female labor across industries, it can be argued that insofar as jobs in manufacturing require relatively more "brawn" than "brain" skills than jobs in services, and women are under-represented in brawn-intensive jobs, the elasticity of substitution between male and female labor may be expected to be higher in service industries than in manufacturing (see also the discussion by Rendall, 2010). In addition, one can argue that the brain versus brawn distinction would matter mostly for the unskilled, as supposedly unskilled jobs in services and manufacturing would involve quite different combinations of brain and brawn, but not so much for the skilled, who are more likely to hold managerial, professional or administrative positions in either sector. We thus consider two cases. In row 3 we assume $\sigma_{serv} = 3$ and $\sigma_{goods} = 1.5$, which, given an average weight of services across countries around 70%, imply an average elasticity of about 2.5, coinciding with our benchmark. In this case the between-industry component of labor demand explains just about the whole variation (103.8%). In row 4 we allow σ to vary across both skills and sectors. Specifically, we set $\sigma_{U,serv} = 3$, $\sigma_{U,goods} = 1.5$, and $\sigma_{S,serv} = \sigma_{S,goods} = 2.5$, and obtain a weight of the between-industry component about 80%.

In summary, in the simplest, benchmark case with constant σ across skills and sectors, we compute that the between-industry component explain as much as 80% of overall differences in labor demand. We then let σ vary within a plausible range defined by available estimates, and finally let it differ across industries and skills, and still identify a substantial impact of the between-industry component in shaping international differences in labor demand. The most flexible, and possibly most realistic, scenario, in which low-skill men and women are more substitutable in services than in manufacturing delivers an estimate of between-industry forces that is remarkably close to that obtained in our benchmark example.

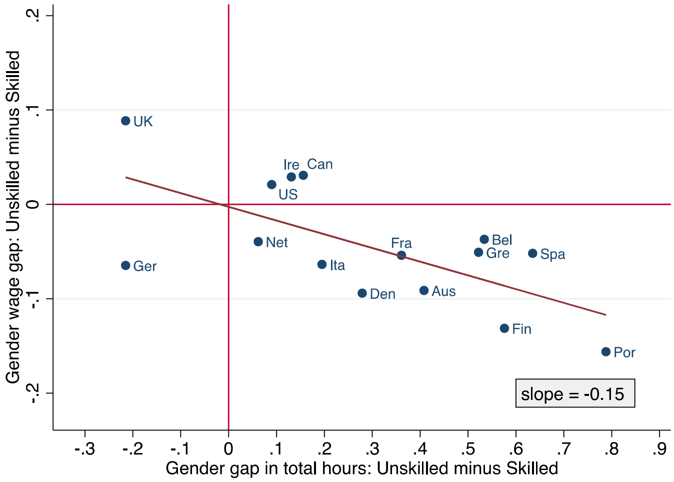

Notes. Adjusted wage gaps are obtained from Eq. (20) in the text for $\sigma = 2.5$ and $\eta = 0$. Hours gaps are gender differences in log hours by country and skill, using population weights.

Fig. 4. Gender gaps in wages and total hours: Unskilled-to-skilled differences adjusted for the between-industry component of labor demand.

4.2. The between-industry component of gender gaps

To quantitatively assess the role of the between-industry component of labor demand on gender outcomes, this section shows two simple quantitative exercises. The first exercise obtains the counterfactual correlation between the unskilled-to-skilled wage gap and the corresponding hours gap that would be observed in equilibrium, having corrected wage and hours gaps for the between-industry component of labor demand. According to the expressions for equilibrium gender gaps (11) and (12), adjusted triple differences in wages and hours are given by:

$$\Delta_C(\Delta w_U - \Delta w_S) - \frac{\sigma}{\sigma+\eta}\left(\Delta_C\tilde{\beta}_U^{between} - \Delta_C\tilde{\beta}_S^{between}\right) \tag{20}$$

and

$$\Delta_C(\Delta h_U - \Delta h_S) - \frac{\sigma\eta}{\sigma+\eta}\left(\Delta_C\tilde{\beta}_U^{between} - \Delta_C\tilde{\beta}_S^{between}\right) \tag{21}$$

respectively, where the $\Delta_C\tilde{\beta}_i^{between}$ terms are given by Eq. (15).[12]

In order to evaluate (20) and (21), we need estimates for both labor demand and labor supply elasticities. For the labor demand elasticity we use our benchmark value of $\sigma = 2.5$. As for the elasticity of labor supply, the consensus in the micro literature is that labor supply elasticities are fairly small, certainly below 1 (see Blundell and Macurdy, 1999). However, Keane and Rogerson (2012) show that relatively small elasticities at the individual level can be consistent with larger elasticities at the aggregate level, in a range between 1 and 2, as typically assumed in general equilibrium models.[13]

In the limiting case $\eta = 0$, hours gaps would be unaffected by demand differences, and adjusted wage gaps would be simply $\Delta_C(\Delta w_U - \Delta w_S) - \Delta_C(\tilde{\beta}_U^{between} - \tilde{\beta}_S^{between})$. We plot these adjusted wage gaps against hours gaps in Fig. 4, and note that the slope of the new fitted line equals -0.15. Recall that the corresponding relationship between actual gaps was essentially flat (Fig. 2), while the theoretical relationship that one should obtain in the absence of relative demand differences would have slope $-1/\sigma = -0.4$. This implies that between-industry demand differences would absorb 37.5% (0.15/0.4) of the observed comovement in wage and hours gaps. This proportion falls slightly to 31% if $\eta = 1$ (and wage *and* hours gaps are adjusted accordingly). The relative importance of the between-industry component declines very little thereafter, falling to 28% for $\eta = 2$, and to 26% for $\eta = 8$. The importance of the between-industry component is thus not very sensitive to the assumed value of the labor supply elasticity, and it remains quantitatively important even when η is set at a very generous upper bound. We would thus conclude that for a plausible range of η estimates, the between-industry

[12] Incidentally, this simple framework is no longer valid if supply elasticities are gender-specific. However, it can be shown that, under plausible assumptions, all that matters for measuring the impact of the between-industry component on the gender wage gap is the average elasticity of labor supply across genders. In the simulations below we are considering very generous upper bounds for such elasticity, providing lower bounds for the impact of the between-industry components on wage gaps. In general, there is only scant evidence on gender differences in the *macro* elasticity of labor supply, and the few studies that consider a gender dimension do not suggest systematic difference between men and women (see for example Chetty, 2012, Table 1).

[13] The micro/macro elasticity puzzle has generated an active strand of work in the labor supply literature in recent years. The channels driving the observed gap between micro and macro elasticities include the social multiplier (Glaeser et al., 2003; Maurin and Moschion, 2009), extensive margin adjustments (Rogerson and Wallenius, 2009), and optimization frictions (Chetty et al., 2011; Chetty, 2012).

Table 4
Unskilled-to-skilled difference in gender gaps across countries: the importance of services.

	Dependent variable: Unskilled-to-skilled wage gap				
	(1)	(2)	(3)	(4)	(5)
Slope	−0.019	−0.204**	−0.031	−0.033	−0.055
(p-value)	(0.684)	(0.043)	(0.674)	(0.514)	(0.387)
R-squared	0.007	0.802	0.153	0.109	0.250
Observations	15	15	15	15	14
		Wage bill shares in			
Other controls	–	Six service industries	Primary; manuf; constr	Public sector	ICT share

Notes. The table reports coefficients and p-values from regressions of the unskilled-to-skilled wage gap on the unskilled-to-skilled hours gap, controlling for other factors in turn. See notes to Table 1 and Fig. 2 for definitions, sample and source. *Other controls*. Columns 2 to 4: Wage bill shares obtained on our main sample (six extra regressors in column 2, three extra regressors in column 3, one extra regressor in column 4). The six service industries are transport, storage, and post and telecommunications; wholesale and retail trade, and hotels and restaurants; financial intermediation, insurance and real estate; education; health and social work; other services. The primary sector includes agriculture, hunting, forestry, fishing, mining, quarrying, electricity, gas and water supply. The public sector includes public administration and defence. Column 5: IT capital share in total capital compensation, obtained from EU Klems, March 2008 release (available at http://www.euklems.net/). Data for Greece are not available.

component of labor demand differences explains close to one third of the observed correlation between wage and hours gaps.

The second exercise consists in estimating the slope coefficient between wage and hours gaps, having controlled for various indicators of the industry structure. The raw correlation between the two variables is close to zero, and we would expect it to turn negative and significant whenever demand differences are properly controlled for. The results of these simple regressions are reported in Table 4. In column 1 we regress the unskilled-to-skilled wage gap on the corresponding hours gap, without any control for the industry structure, and the results simply replicate the flat relationship shown graphically in Fig. 2. Column 2 includes controls for wage bill shares in six service industries, and the slope turns negative and significant. In particular, a slope of −0.122 is equal to about one third of the theoretical slope that one would observe in the presence of pure labor supply differences (−0.4). Thus the variation in the share of service industries explains as much as the whole between industry component – as observed in the previous exercise – implying that the bulk of between-industry forces lies in the cross-country variation of the service share. To further illustrate this point, columns 3 and 4 control for the shares of goods-producing industries and the public sector share, respectively, and in each case the relationship between wage and hours gap remains flat. Finally, column 5 controls for the IT share in total capital compensation, which may affect demand for skills and genders *within* industries. In this case the slope turns slightly more negative than in column 1 but it is not statistically significant. Of course, given the basic specification and small sample size, one should take these correlations with more than some caution. With this qualification in mind, the figures reported in Table 4 suggest that the share of service industries can absorb a sizable portion of the cross-country variation in wage and hours gaps.

4.3. Possible explanations

An explanation often discussed for the variation in the services share across countries relates to differences in the rate of marketization of home production. When household activities like childcare, elderly care, cooking, house repairs, gardening, etc. are outsourced to the market, they accrue to the broad service sector. Freeman and Schettkat (2005) provide evidence on the marketization hypothesis, based on both time-use data and expenditure data across countries, and conclude that it contributes substantially to the hours gap across the Atlantic. In a similar vein, Rogerson (2008) relates the relative poor performance of continental European labor markets to an under-marketized service industry. Marketization of services may in turn be hindered in continental Europe by a higher tax wedge, which distorts market-home substitution. Ngai and Pissarides (2011) provide evidence on this mechanism for a number of OECD countries by showing that taxation and subsidies decrease and raise hours, respectively, in sectors that have close home substitutes. Finally, as women are the primary provider of home services, social norms about women's work in and out of the household may also have an impact on the service share.

Another potential driving force for the rise of services is the productivity growth differential between manufacturing and service industries. As manufacturing output and services are poor substitutes in consumption, faster productivity growth in manufacturing reallocates labor from manufacturing into services. While this hypothesis is typically framed in a historical perspective to explain the secular rise in services (see, among others, Ngai and Pissarides, 2008, and Herrendorf et al., 2014), similar intuition implies that international differences in productivity growth would map out into differential growth in the service share in a cross-section of countries.

Below we explore country characteristics that are potentially related to the share of services, in line with the hypotheses discussed above. In particular we consider a number of institutional factors and cultural indicators that may directly affect the marketization of home production, and the productivity growth differential between goods-producing industries and service industries.

Table 5
The service share: correlation with institutions, attitudes, and productivity.

Dependent variable	Service hours share								Change in hours share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Regressor	Marginal tax rate 2nd earner	Average tax wedge	Public spending childcare, pre-school	EPL	Attitudes about gender roles				Productivity growth: goods-services
					Scarce jobs (Women)	Scarce jobs (Men)	Housewife (Women)	Housewife (Men)	
Slope	−0.023**	0.0178	0.048*	−0.0258**	−0.280**	−0.085	0.134	0.188	0.0353
(p-value)	(0.013)	(0.862)	(0.056)	(0.017)	(0.035)	(0.426)	(0.252)	(0.115)	(0.560)
R-squared	0.167	0.001	0.114	0.269	0.196	0.061	0.094	0.143	0.058
Observations	15	15	15	14	15	15	15	15	15

Notes. Columns 1 to 8 report coefficients and associated p-values obtained by regressing the average hours share in the service sector on various country-level indicators. *Regressors included*. Column 1: Marginal income tax rate for the spouse of a two-earner married family with 2 children in which the head earns 100% of the average gross wage (APW) and the spouse earns 67% of the APW, divided by the corresponding single earner marginal tax rate. Source: Column 6, Table 7, in OECD *Taxing Wages* 2000–2001 (available at http://www.oecd.org/tax/taxingwages.htm, 2000). Column 2: Average tax wedge computed by adding (a) the average tax rates from OECD *Taxing Wages* statistics including employers' social security contributions and (b) the average consumption tax rate ((Indirect taxes − subsidies) ÷ consumption) from OECD National Accounts, 2000; Columns 3: Public spending in childcare and pre-school as a percentage of GDP, average 1998–2001; Sources: Table PF2.1.A and PF3.1, in OECD (2012); Column 4: Employment protection legislation indicator for regular work, 2000. Source: Nickell (2006); data for Greece not available. Columns 5 and 6: Attitudes toward gender roles, measured as mean response in World Value Survey to the statement "*When jobs are scarce, men should have more right to a job than women*" (0–1 scale: 0 indicates no agreement with the statement, 1 indicates complete agreement with the statement). Women's views in column 5; Men's views in column 6. Columns 7 and 8: Attitudes toward gender roles measured as mean response in World Value Survey to the statement "*Being a housewife is just as fulfilling as working for pay*" (0–1 scale: 0 indicates no agreement with the statement, 1 indicates complete agreement with the statement). Women's views in column 7; Men's views in column 8. Source for columns 5–8: Fortin (2005, Appendix Table 2, columns 1, 3, 9, 11). Sample period: average over 1990–1993, 1995–1997, and 1999–2001. Column 9 regresses the average change in the service hours share over the sample period on the average differential in real productivity growth between goods-producting industries (primary, utilities, manufacturing and construction) and service-producing industries (transport and storage; post and telecommunications; wholesale and retail trade; hotels and restaurants; financial intermediation, insurance and real estate; education; health and social work; other services; public administration and defense.) Source: EU Klems, March 2011 update of November 2009 release (March 2008 release for Canada), available at http://www.euklems.net/. Reference years 1993–2001 (1996–2004 for Canada).

Columns 1–4 in Table 5 show slope coefficients between the service hours share and a few institutional variables. A higher tax rate for secondary earners is associated with a smaller service share, while public provision of childcare is associated with a larger service share, and both effects are in line with the marketization hypothesis. The strictness of EPL is another element of labor regulation that is negatively related to the size of services. Columns 5–8 focus on social attitudes towards female work, and show that the proportion of women believing than men are more deserving of work in bad times is associated to a smaller service share (but slope coefficients for other indicators of attitudes are not statistically significant). Finally column 9 regresses the change in the service share on the inter-industry productivity growth differential, and the obtained coefficient is positive, but it is not statistically significant. While column 9 solely exploits cross-country variation for consistency with the rest of the table, we also estimate a specification similar to 9 on a country panel for all years available, controlling for country fixed effects, and thus relying on within-country variation. In this case we obtain a coefficient on the productivity growth differential of 0.12, which is significant at the 5% level, and consistent with the hypothesis that uneven productivity growth is an important predictor of the rise of services.

5. Conclusions

This paper uncovers a strong, positive correlation between the unskilled-to-skilled wage gap and the corresponding gap in hours per head across countries, thus pointing at significant (net) demand forces shaping gender differences in labor market outcomes across skills. We link the structure of labor demand to cross-country differences in the process of structural transformation, and specifically in the weight of services. Based on a stylized, multi-sector model of labor demand, we decompose the gender bias in labor demand into measurable between-industry and within-industry components. Using comparable micro data across countries, we conclude that differences in the industry structure explain over 80% of overall labor demand differences, and about one third of the correlation between wage and hours gaps.

Appendix A. Derivation of model results

A.1. Derivation of (6), (7) and (15)

The F.O.C. for skilled male wages is

$$W_{MSj} = P_j \frac{\partial Q_j}{\partial M_{ij}} = \theta_j Q_j^{-1}\left[\alpha_j S_j^{\frac{\phi-1}{\phi}} + (1-\alpha_j) U_j^{\frac{\phi-1}{\phi}}\right]^{\frac{1}{\phi-1}} \alpha_j S_j^{-\frac{1}{\phi}}$$
$$\times \left[\alpha_{S_j}(B_{MSj} M_{Sj})^{\frac{\sigma-1}{\sigma}} + (1-\alpha_{S_j})(B_{FSj} F_{Sj})^{\frac{\sigma-1}{\sigma}}\right]^{\frac{1}{\sigma-1}} \alpha_{S_j} B_{MSj}^{\frac{\sigma-1}{\sigma}} M_{Sj}^{-\frac{1}{\sigma}},$$

having normalized $Q_r = 1$. Rearranging, we obtain

$$W_{MSj} = \theta_j \tilde{\alpha}_j \beta_{S_j} B_{MSj}^{\frac{\sigma-1}{\sigma}} S_j^{\frac{1-\sigma}{\sigma}} M_{Sj}^{-\frac{1}{\sigma}}, \tag{22}$$

where

$$\tilde{\alpha}_j = \frac{\alpha_j S_j^{\frac{\phi-1}{\phi}}}{\alpha_j S_j^{\frac{\phi-1}{\phi}} + (1-\alpha_j) U_j^{\frac{\phi-1}{\phi}}}$$

represents the wage bill share of skilled labor.

Perfect labor mobility implies wage equalization, i.e. $W_{MSj} = W_{MS}$ for all j. Using this property and adding (22) across sectors gives

$$\sum_j M_{Sj} = M_S = \frac{\sum_j \theta_j^\sigma \tilde{\alpha}_j^\sigma \beta_{S_j}^\sigma B_{MSj}^{\sigma-1} S_j^{1-\sigma}}{W_{MS}^\sigma}.$$

Solving for W_{MS} yields:

$$W_{MS} = \left(\sum_j \theta_j^\sigma \tilde{\alpha}_j^\sigma \beta_{S_j}^\sigma B_{MSj}^{\sigma-1} S_j^{1-\sigma}\right)^{\frac{1}{\sigma}} M_S^{-\frac{1}{\sigma}}. \tag{23}$$

Combining (23) with the corresponding expression for skilled female wages finally gives $\Delta w_S = \tilde{\beta}_S - \frac{1}{\sigma}\Delta h_S$, which is equivalent to (5), with $\tilde{\beta}_S$ given by expression (6). Similar steps for W_{MU} yield expression (7).

One can next differentiate $\tilde{\beta}_S$ (respectively, $\tilde{\beta}_U$) with respect to θ_j to obtain the between-industry component of cross-country differences in labor demand for the skilled (respectively, the unskilled):

$$\Delta_C \tilde{\beta}_S^{between} = \frac{\sum_j x_{MSj,C} \theta_{j,C}^{\sigma-1} \Delta_C \theta_j}{\sum_j x_{MSj,C} \theta_{j,C}^{\sigma}} - \frac{\sum_j x_{FSj,C} \theta_{j,C}^{\sigma-1} \Delta_C \theta_j}{\sum_j x_{FSj,C} \theta_{j,C}^{\sigma}}, \tag{24}$$

where

$$x_{MSj} \equiv \tilde{\alpha}_j^{\sigma} \beta_{S_j}^{\sigma} B_{MSj}^{\sigma-1} S_j^{1-\sigma}$$

$$x_{FSj} \equiv \tilde{\alpha}_j^{\sigma} (1-\beta_{Sj})^{\sigma} B_{FSj}^{\sigma-1} S_j^{1-\sigma}.$$

Using (22) and the corresponding expression for skilled females, x_{MSj} and x_{FSj} can be rewritten as

$$x_{MSj} = W_{MS}^{\sigma} M_{Sj} \theta_j^{-\sigma} \tag{25}$$

$$x_{FSj} = W_{FS}^{\sigma} F_{Sj} \theta_j^{-\sigma}. \tag{26}$$

Substituting (25) and (26) into (24) gives:

$$\Delta_C \tilde{\beta}_S^{between} = \frac{\sum_j y_{MSj,C} \Delta_C y_{j,C}}{y_{MS,C}} - \frac{\sum_j y_{FSj,C} \Delta_C y_{j,C}}{y_{FS,C}}, \tag{27}$$

where y_j denotes industry j's share of the total wage bill, y_{Mij} denotes the wage bill share of males of skill i in industry j, and y_{Mi} denotes the wage bill share of males of skill i in the economy. Note that wage equalization implies that the $W_{MS}^{\sigma-1}$ and $W_{FS}^{\sigma-1}$ terms cancel out from numerators and denominators when we substitute (25) and (26) into Eq. (24), leading to (27).

The differentiation in (24) and thus in (27) would be exact (and unique) for infinitesimal changes in the underlying variables, while we are considering discrete changes between country C and the U.S. In other words, when differentiating (6) one may derive equivalent expressions to (24) and (27) in which all level variables refer to the U.S. as opposed to country C. To limit arbitrariness, in the empirical evaluation we express all level variables in (27) as simple averages across country C and the U.S. This gives expression (15) for $i = S$, and similar steps can be repeated to obtain $\Delta_C \tilde{\beta}_U^{between}$.

A.2. Derivation of (19)

Let's denote by σ_{ij} the elasticity of substitution between male and female labor for skill group i and industry j. The F.O.C. for skilled male wages is given by

$$W_{MSj} = \theta_j \tilde{\alpha}_j \beta_{S_j} B_{MSj}^{\frac{\sigma_{Sj}-1}{\sigma_{Sj}}} S_j^{\frac{1-\sigma_{Sj}}{\sigma_{Sj}}} M_{Sj}^{-\frac{1}{\sigma_{Sj}}}. \tag{28}$$

Imposing wage equalization and adding across industries gives:

$$\sum_j M_{Sj} = M_S = \sum_j \theta_j^{\sigma_{Sj}} \tilde{\alpha}_j^{\sigma_{Sj}} \beta_{S_j}^{\sigma_{Sj}} B_{MSj}^{\sigma_{Sj}-1} S_j^{1-\sigma_{Sj}} W_{MS}^{-\sigma_{Sj}}. \tag{29}$$

Next define the cross-industry average of the elasticity of substitution for the skilled, $\bar{\sigma}_S$. This allows us to rewrite (29) as

$$M_S = W_{MS}^{-\bar{\sigma}_S} \sum_j \theta_j^{\sigma_{Sj}} x_{SjM} W_{MS}^{\bar{\sigma}_S - \sigma_{Sj}}. \tag{30}$$

Combining (30) with the corresponding expression for skilled females gives $\Delta w_S = \tilde{\beta}_S - \frac{1}{\bar{\sigma}_S} \Delta h_S$, with $\tilde{\beta}_S$ given by expression (18). Similar steps for W_{MU} yield an expression equivalent to (18) for the unskilled.

Differentiating $\tilde{\beta}_S$ (respectively, $\tilde{\beta}_U$) with respect to θ_j gives the between-industry component of cross-country differences in labor demand for the skilled (respectively, the unskilled):

$$\Delta_C \tilde{\beta}_S^{between} = \frac{1}{\bar{\sigma}_S} \left[\frac{\sum_j x_{MSj,C} \sigma_{Sj} W_{MS,C}^{(\bar{\sigma}_S - \sigma_{Sj})} \theta_{j,C}^{\sigma_{Sj}-1} \Delta_C \theta_j}{\sum_j x_{MSj,C} W_{MS,C}^{(\bar{\sigma}_S - \sigma_{Sj})} \theta_{j,C}^{\sigma_{Sj}}} - \frac{\sum_j x_{FSj,C} \sigma_{Sj} W_{FS,C}^{(\bar{\sigma}_S - \sigma_{Sj})} \theta_{j,C}^{\sigma_{Sj}-1} \Delta_C \theta_j}{\sum_j x_{FSj,C} W_{FS,C}^{(\bar{\sigma}_S - \sigma_{Sj})} \theta_{j,C}^{\sigma_{Sj}}} \right], \tag{31}$$

where

$$x_{MSj} \equiv \tilde{\alpha}_j^{\sigma_{Sj}} \beta_{S_j}^{\sigma_{Sj}} B_{MSj}^{\sigma_{Sj}-1} S_j^{1-\sigma_{Sj}}$$

$$x_{FSj} \equiv \tilde{\alpha}_j^{\sigma_{Sj}} (1-\beta_{Sj})^{\sigma_{Sj}} B_{FSj}^{\sigma_{Sj}-1} S_j^{1-\sigma_{Sj}}.$$

Using (28) these can be rewritten as:

$$x_{MSj} = W_{MS}^{\sigma_{Sj}} M_{Sj} \theta_j^{-\sigma_{Sj}}; \tag{32}$$

$$x_{FSj} = W_{FS}^{\sigma_{Sj}} F_{Sj} \theta_j^{-\sigma_{Sj}}. \tag{33}$$

Substituting (32) and (33) into (31) and using again cross-country averages for level variables gives (19) for $i = S$. Similar steps can be repeated to obtain $\Delta_C \tilde{\beta}_U^{between}$.

Appendix B. Descriptive statistics on the sample used

Table B1
Distribution of population by education.

	Males			Females		
	Educ. = 1	Educ. = 2	Educ. = 3	Educ. = 1	Educ. = 2	Educ. = 3
U.S.	13.8	50.1	36.1	12.4	51.3	36.4
Canada	20.4	63.4	16.3	17.6	65.6	16.8
U.K.	37.4	19.7	43.0	49.6	18.9	31.5
Finland	19.8	46.4	33.8	17.4	36.2	46.4
Denmark	17.3	43.9	38.8	18.6	37.6	43.8
Germany	17.5	57.2	25.3	23.2	59.0	17.8
Netherlands	24.1	52.4	23.5	29.4	52.2	18.4
Belgium	26.6	34.7	38.7	26.6	31.8	41.6
Austria	12.4	79.3	8.3	26.2	64.1	9.6
Ireland	44.5	35.2	20.3	42.9	41.5	15.6
France	29.7	44.9	25.4	33.9	38.3	27.8
Italy	47.4	41.9	10.7	51.1	39.8	9.1
Spain	56.0	19.1	24.9	59.9	16.9	23.1
Portugal	79.3	13.7	7.0	77.4	12.5	10.0
Greece	36.2	35.2	28.6	46.3	28.4	25.3

Notes. Educ. = 1 includes individuals with less than upper secondary education; Educ. = 2 includes individuals who have completed upper secondary education; Educ. = 3 includes individuals who have completed college education or above. See notes to Table 1 for sample and source.

Table B2
Wage bill shares of four demographic groups and gender differences.

	Unskilled (no college degree)			Skilled (college degree)		
	Males (1)	Females (2)	Difference (1)–(2)	Males (1)	Females (2)	Difference (1)–(2)
U.S.	29.73	18.45	11.29	31.09	20.73	10.36
Canada	45.16	29.38	15.77	13.91	11.55	2.37
U.K.	28.44	18.87	9.58	34.06	18.62	15.44
Finland	31.66	19.24	12.42	24.29	24.82	−0.53
Denmark	30.48	20.33	10.14	26.81	22.38	4.43
Germany	42.96	24.95	18.00	22.28	9.80	12.48
Netherlands	46.18	21.83	24.35	21.90	10.09	11.81
Belgium	30.82	15.02	15.80	29.90	24.26	5.64
Austria	58.18	27.76	30.42	8.44	5.62	2.82
Ireland	43.16	21.38	21.78	21.75	13.71	8.05
France	37.78	20.93	16.85	23.99	17.29	6.70
Italy	54.77	28.67	26.09	10.12	6.44	3.68
Spain	41.43	14.27	27.15	27.24	17.05	10.19
Portugal	49.13	26.77	22.35	10.97	13.13	−2.17
Greece	40.79	16.12	24.67	25.43	17.65	7.78

See notes to Table 1 for sample and source.

References

Acemoglu, Daron, Autor, David H., 2011. Skills, tasks and technologies: implications for employment and earnings. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 4B. Elsevier–North Holland, Amsterdam, pp. 1043–1171. Ch. 12.
Acemoglu, Daron, Autor, David H., Lyle, David, 2004. Women, war, and wages: the effect of female labor supply on the wage structure at mid-century. Journal of Political Economy 112 (3), 497–551.
Akbulut, Rahşan, 2011. Sectoral changes and the increase in women's labor force participation. Macroeconomic Dynamics 15, 240–264.
Autor, David H., Dorn, David, 2013. The growth of low-skill service jobs and the polarization of the U.S. labor market. The American Economic Review 103, 1553–1597.
Azmat, Ghazala, Güell, Maia, Manning, Alan, 2006. Gender gaps in unemployment rates in OECD countries. Journal of Labor Economics 24 (1), 1–38.

Black, Sandra E., Spitz-Oener, Alexandra, 2010. Explaining women's success: technological change and the skill content of women's work. Review of Economics and Statistics 92 (1), 187–194.
Blau, Francine D., Kahn, Lawrence M., 1996. Wage structure and gender earnings differentials: an international comparison. Economica 63, 29–62.
Blau, Francine D., Kahn, Lawrence M., 2003. Understanding international differences in the gender pay gap. Journal of Labor Economics 21, 106–144.
Blundell, Richard, Macurdy, Thomas, 1999. Labor supply: a review of alternative approaches. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3. Elsevier, pp. 1559–1695.
Bound, John, Johnson, George, 1992. Changes in the structure of wages in the 1980s: an evaluation of alternative explanations. The American Economic Review 82, 371–392.
Chetty, Raj, 2012. Bounds on elasticities with optimization frictions: a synthesis of micro and macro evidence on labor supply. Econometrica 80 (3), 969–1018.
Chetty, Raj, Friedman, John N., Olsen, Tore, Pistaferri, Luigi, 2011. Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: evidence from Danish tax records. The Quarterly Journal of Economics 126 (2), 749–804.
Fortin, Nicole M., 2005. Gender role attitudes and the labour market outcomes of women across OECD countries. Oxford Review of Economic Policy 21, 416–438.
Freeman, Richard B., Schettkat, Ronald, 2005. Marketization of household production and the EU-US gap in work. Economic Policy 20, 6–50.
Fuchs, Victor R., 1968. The Service Economy. National Bureau of Economic Research, Inc.
Glaeser, Edward L., Sacerdote, Bruce I., Scheinkman, Jose A., 2003. The social multiplier. Journal of the European Economic Association 1 (2–3), 345–353.
Goldin, Claudia, 1995. The U-shaped Female labor force function in economic development and economic history. In: Schultz, T.P. (Ed.), Investment in Women's Human Capital and Economic Development, pp. 61–90.
Goldin, Claudia, 2006. "The quiet revolution that transformed women's employment, education, and family". Ely lecture. The American Economic Review, Papers and Proceedings 96, 1–21.
Goos, Maarten, Manning, Alan, Salomons, Anna, forthcoming. The polarization of the European labor market. The American Economic Review.
Hamermesh, Daniel S., 1993. Labor Demand. Princeton University Press, Princeton, NJ.
Heathcote, Jonathan, Storesletten, Kjetil, Violante, Gianluca, 2010. The macroeconomic implications of rising wage inequality in the United States. Journal of Political Economy 118 (4), 681–722.
Herrendorf, Berthold, Rogerson, Richard, Valentinyi, Àkos, 2014. Growth and structural transformation. In: Aghion, P., Durlauf, S. (Eds.), Handbook of Economic Growth, vol. 2, pp. 855–941. Ch. 6.
Johnson, Matthew, Keane, Michael P., 2013. A dynamic equilibrium model of the US wage structure, 1968–1996. Journal of Labor Economics 31 (1), 1–49.
Katz, Lawrence F., Autor, David H., 1999. Changes in the wage structure and earnings inequality. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3A. Elsevier Science, Amsterdam, pp. 1463–1555.
Katz, Lawrence F., Murphy, Kevin M., 1992. Changes in relative wages, 1963–1987: supply and demand factors. The Quarterly Journal of Economics 107 (1), 35–78.
Keane, Michael P., Rogerson, Richard, 2012. Micro and macro labor supply elasticities: a reassessment of conventional wisdom. Journal of Economic Literature 50 (2), 464–476.
King, Miriam, Ruggles, Steven, Alexander, J. Trent, Flood, Sarah, Genadek, Katie, Schroeder, Matthew B., Trampe, Brandon, Vick, Rebecca, 2010. Integrated public use microdata series, current population survey: version 3.0 [Machine-readable database]. University of Minnesota, Minneapolis.
Lee, Donghoon, Wolpin, Ken, 2006. Intersectoral labor mobility and the growth of the service sector. Econometrica 74, 1–46.
Maurin, Eric, Moschion, Julie, 2009. The social multiplier and the labor market participation of mothers. American Economic Journal: Applied Economics 1 (1), 251–272.
Nickell, W., 2006. The CEP–OECD institutions data set. CEP Discussion Paper No. 759.
Ngai, L. Rachel, Pissarides, Christopher A., 2008. Trends in hours and economic growth. Review of Economic Dynamics 11 (2), 239–256.
Ngai, L. Rachel, Pissarides, Christopher A., 2011. Taxes, social subsidies and the allocation of work time. American Economic Journal: Macroeconomics 3, 1–26.
Ngai, Rachel, L., Petrongolo, Barbara, 2014. Gender gaps and the rise of the service economy. Unpublished.
OECD, 2012. OECD family database. OECD, Paris. http://www.oecd.org/social/family/database.
Olivetti, Claudia, 2014. The female labor force and long-run development: the American experience in comparative perspective. In: Platt Boustan, L., Frydman, C., Margo, R.A. (Eds.), Human Capital in History: The American Record. University of Chicago Press.
Olivetti, Claudia, Petrongolo, Barbara, 2008. Unequal pay or unequal employment? A cross-country analysis of gender gaps. Journal of Labor Economics 26, 621–654.
Rendall, Michelle, 2010. Brain versus Brawn: the realization of women's comparative advantage. Unpublished.
Rendall, Michelle, 2013. The service sector and female market work: Europe vs. U.S. Unpublished.
Rogerson, Richard, 2007. Taxation and market work: is Scandinavia an outlier?. Economic Theory 32, 59–85.
Rogerson, Richard, 2008. Structural transformation and the deterioration of European labor markets. Journal of Political Economy 166, 235–258.
Rogerson, Richard, Wallenius, Johanna, 2009. Micro and macro elasticities in a life cycle model with taxes. Journal of Economic Theory 144, 2277–2292.
Weinberg, Bruce A., 2000. Computer use and the demand for female workers. Industrial & Labor Relations Review 53, 290–308.