THE RIMINI CENTRE FOR ECONOMIC ANALYSIS

WP 15-42

Manuel Hidalgo-Perez
Universidad Pablo de Olavide, Spain

Benedetto Molinari
Universidad Pablo de Olavide, Spain
The Rimini Centre for Economic Analysis, Italy

LEARNING NEW TECHNOLOGY: THE POLARIZATION OF THE WAGE DISTRIBUTION

Learning New Technology: The Polarization of the Wage Distribution*

Manuel Hidalgo-Pérez
Universidad Pablo de Olavide

Benedetto Molinari
Universidad Pablo Olavide and RCEA

November 12, 2015

Abstract

This paper revisits the relationship between wage inequality and technological progress. By applying counterfactual quantile regressions to historic U.S. data, we show that the reduction of wage inequality among low-wage workers generated by routinization-biased technical change was fully driven by a reduction of within-group inequality, which was determined by more homogeneous remunerations paid to routine workers. Changes in wage differentials between workers performing technology-neutral and technology-substitute tasks played instead a negligible role, which casts some doubt on a theory of technical change operating through a labor-demand channel. To reconcile the theory with data, we develop a model in which skill-heterogeneous workers face endogenous occupational choices and learning costs in connection with operating a new technology. In this model, when wage differentials are fixed technical change still generates an empirically-consistent non-monotone effect on wage inequality by affecting the average levels of skills within different groups of workers.

JEL Classification: J24, J31, O33

Keywords: Residual Wage Inequality, Wage Polarization, Price and Composition Effect, Routinization hypothesis, Skill Biased Technical Change, Occupational Tasks.

*We wish to acknowledge useful comments on this paper from David Dorn, Antonio Villar, Nacho Garcia-Pérez, the participants at the CEMFI Macro Workshop, the 28th EEA Annual Meeting, the XXXIX SAEe Annual Meeting, the X AEET Meeting, and the XXIX AIEL National Conference. Corresponding author: Manuel Hidalgo-Peréz, *mhidper@upo.es*, Department of Economics, UPO, Sevilla (Spain - EU)

1 Introduction

During the postwar period, wage inequality in the U.S. remained relatively stable until the end of the 1970s, when it began to rise noticeably [see Juhn, Murphy and Pierce (1993) for wage inequality during the 1980s, Acemoglu (2002) and Lemieux (2006) during the 1990s, and Acemoglu and Autor (2010) during the 2000s.] This evidence stimulated a substantial debate about the concurrent causes that might have led to such an increase, including institutional factors such as declining minimum wages and de-unionization [Freeman and Katz (1995); Di Nardo, Fortin and Lemieux (1996)], greater commercial openness and trade [Acemoglu (2003)], and technological progress biased toward skilled employment [Juhn, Murphy and Pierce (1993)].[1] Although the latter became the mainstream explanation in the literature, several empirical studies questioned its role as key determinant of wage inequality by highlighting several aspects of labor market data that appear to be inconsistent with the theory or that the theory is unable to rationalize [Freeman and Katz (1995), Buchinsky (1998), DiNardo, Fortin and Lemieux (1996), Piketty and Saez (2003, 2006), Lemieux (2006)]. In response to such criticism, some authors suggested a more nuanced version of skill-biased technical change to reappraise the relationship between technology and wage inequality. By studying the effects of computers, ICT and automated machines on business practices, Autor, Levy and Murnane (2003) argued that technical change is not biased toward all skills, but is complementary to some of them and substitute to others.[2] As a result, technical change is expected to raise only the wages of workers performing technology-complementary tasks while reducing the wages of workers performing technology-substitute tasks, rather than generally increasing skill premia.[3]

[1] Skill-biased technical change is intended as the effects of continuously growing technological progress – such as that generated by computers, ICT and electronically controlled machines – on the demand for skilled workers who are capable of operating such new technologies. By stimulating the demand for skilled labor, technology increases the remuneration of skilled workers relative to unskilled workers (the skill premium), thus pushing wage dispersion upward. Technical change has been shown to be biased towards skills by Bound and Johnson (1992), Katz and Murphy (1992), Levy and Murnane (1992), Krueger (1993), Berman et al. (1994), among others.

[2] For instance, in manufacturing plants technological progress replaced traditional "blue collar" workers with automated machines, which enhanced the productivity of "white collar" officers in charge of the assembly lines. Thus, it created a comparative advantage of cognitive skills with respect to manual dexterity skills that eventually increased the wages of white collar workers relative to blue collar workers.

[3] See Acemoglu and Autor (2011) and references therein. In this literature, a task is defined as a unit of working activity that produces output, and a skill is defined as a worker's ability to perform a designated task. Acemoglu and Autor (2011) provide an exhaustive analysis of the empirical facts that the canonical skill-biased technical change model cannot explain, but that a task-based model can.

Moreover, because occupational tasks are *not* randomly distributed across the wage distribution – occupations consisting mostly of technology-substitute tasks are typically placed in the middle echelon of the wage distribution, whereas occupations consisting mostly of technology-complementary tasks are typically placed in the upper echelon – then technical change is expected to concurrently depress the remunerations of middle-earners and raise the remunerations of high-earners, which thus explains the increase in wage inequality through a polarization process of the wage distribution.[4]

The empirical literature has not yet established conclusive evidence on the relationship between the routinization hypothesis and wage inequality. By using systematic data from the U.S. labor market and clustering workers according to their education, experience and gender, Autor, Katz and Kearny (2008) provided some supporting evidence. On the one hand, they showed that, when controlling for changes in the composition of the labor force, inequality among workers' salaries tended to vary in opposite directions above and below the median wage since the late 1980s.[5] In particular, residual wage inequality [RWI] increased above and diminished below the median wage thus mimicking the behavior of overall wage inequality shown by Buchinsky (1998). On the other hand, they found that employment diminished in jobs consisting mainly of technology-substitute routine tasks, which are placed in the lower echelon of the skill distribution, while it increased in jobs consisting mainly of technology-complementary cognitive tasks, which are placed in the upper echelon.[6] Hence, they argued that the observed dynamics of RWI was consistent with a theory of job polarization in which routinization-biased technical change [RBTC] compresses the wages of middle-skilled routine workers and enhances the wages of high-skilled cognitive workers. Firpo, Fortin and Lemieux (2011) [FFL] provided a direct estimation of the overall effect of occupational tasks on wage inequality. After controlling for institutional and socio-economic explanatory variables (including minimum wage, de-unionization, age and education), they stated find that, "*once passed through the lens of the routinization hypothesis, the effect of skill-biased technology appears to be the most important determinant of wage inequality during the 1990s.*" However, by analysing the evolution of wages

[4]See Autor, Levy and Murnane (2003), Goos and Manning (2007), Autor, Katz and Kearny (2008), and Autor and Dorn (2013), among others.

[5]RWI is the inequality among the components of wages that are not explained by the observable characteristics of wages. The composition effect refers to the possible effect of occupational tasks on the composition of the labor force. See Lemieux (2006).

[6]Similar evidence has been shown by Goos and Manning (2007) for the U.K, by Goos, Manning and Salomons (2009, 2011) for European countries, and by Acemoglu and Autor (2011) for the U.S.

for 250+ detailed occupations, Mishel, Shierholz and Schmitt (2013) find that the effect on wage inequality of wage differentials between technology-complementary and technology-substitute occupations was too mild to consider RBTC as the main explanation for the observed growth in wage inequality.

In this paper, we propose a unified framework to reconcile and extend previous findings. On the one hand, the paper complements the work of FFL by analyzing the mechanisms through which changes in the remunerations of occupational tasks operate on the wage distribution. On the other hand, it rationalizes the results of Mishel, Shierholz and Schmitt within the routinization hypothesis. The time span analyzed is the 1990s, which is the most representative decade to study the effect of technology on wage inequalities according to FFL. Data on wages are collected using the May/ORG Census Samples database supplemented by the fourth edition (1977) and the revised fourth edition (1991) of the U.S. Department of Labor's Dictionary of Occupational Titles. This database is used to group workers according to the usual socioeconomic characteristics *and* to the tasks performed on duty and then to perform several empirical analyses. First, by employing the reweighing kernel approach of Lemieux (2006) we assess the relative importance of *price* vs *composition effect*. When occupational tasks are maintained constant in the reweighing analysis, the contribution of changes in the composition of the labor force is shown to be smaller than has been found in the previous literature, whereas changes in remunerations (price) for workers' characteristics appear the crucial determinant of the evolution of RWI. Second, we check whether occupational tasks are suitable candidates for explaining wage inequality by performing a set of simple OLS regressions in which the overall change in RWI over the considered decade is regressed on changes in task intensities. The results indicate that (i) occupational tasks have a significantly effect on RWI, (ii) their effects on wages vary along the wage distribution, and (iii) the signs of these effect broadly conform with the theoretical predictions.

Based on previous results, we further analyze the observed changes in wage inequality by pursuing a decomposition approach that controls for changes in the composition of the labor force (composition effect) and then separately identifies the effect of occupational tasks on (i) changes in wage differentials between groups of homogeneous workers (*between-group price effect*), (ii) changes in wages dispersion within groups of homogeneous workers (*within-group price effect*). This decomposition is obtained by performing the extension proposed by Autor, Katz and Kearny (2005) of Machado and Mata (2005) counterfactual quantile regressions [CQR]. Compared with other approaches used in the literature to analyze wage inequality, the CQR features the key property that disentangles the price effect affecting between- and within-group wage inequality, and eventually identifies the contributions of occupational tasks in determining each of the two.[7] When applied to U.S. labor market data over the 1986-2002 period, the CQR analysis reveals that the observed decline of wage inequality among low-earners (below the 30th percentile) is almost entirely explained by the reduction in within-group wage inequality, which is determined, in turn, by the reduction in wage dispersion among routine workers. Both the composition effect and the between-group price effect (wage differentials) computed among narrowly defined groups of low-wage workers − same education, experience *and* tasks performed on duty − appear to have marginal effects on wages dispersion. On the contrary, the dynamics of wages among high-earners (above the 60th percentile) appears to be explained in equal parts by positive price and composition effects and both the between- and within-group price effects are primarily determined by increases in the remuneration for cognitive tasks.

In the last section of the paper, we dwell on the implications of previous results for the routinization hypothesis. Although the estimation outcome for the upper echelon of the wage distribution is consistent with the theory, the results for low-earners appear at odd with the theoretical predictions. If the remuneration of routine skills had diminished, we should have observed a negative effect of routine tasks on between-group wage inequality due to lower wage differentials between workers performing routine tasks and other workers, which in the lower echelon of the wage distribution are mainly manual workers whose initial wages were lower than those of routine workers. Instead, our results show that the effect of routine tasks on wage differentials among low-earners is almost negligible. We argue that this evidence does not necessarily imply a dismissal of the routinization hypothesis, but rather points toward a model of RBTC operating through a *labor-supply* channel. We elaborate upon this intuition by developing a task-based model in which ability-heterogenous individuals decide which occupation to undertake and confront learning costs in adjusting their supply of task-specific labor to new technologies. In this model, when the wage differential between manual and routine occupations is maintained constant – i.e., the labor-demand channel is shut down –, technical change still affects the income distribution

[7] Compared with the RIF-regression approach pursued by FFL, the CQR has the disadvantage that does not allow to identify the overall impact of occupational tasks on wage inequality. In the perspective of our analysis, however, this seems a minor issue because we are training our attention on understanding the channels through which technology operates on wages rather than providing an overall assessment.

by contemporaneously pushing downward within-group inequality among low-earners and upward among high-earners akin to the evidence observed in actual data. Such a dynamics is generated by a sorted migration of skilled workers from routine to manual and abstract occupations. Intuitively, after an increase in technical change, the most capable routine workers find it convenient to put extra effort into obtaining (what is now) a relatively better paid abstract occupation (upward migration), whereas the less-capable routine workers experience an increase in the costs of learning the new technology that makes the choice of maintaining their routine jobs not optimal, and therefore switch to manual occupations (downward migration). Hence, the group of routine workers becomes not only smaller but also more homogeneous in terms of skills and, accordingly, its within-group inequality diminishes as observed in the data. The existence of a sorted migration of routine workers to manual and abstract occupations has been empirically supported by Cortes (2016) and Groes et al. (2015), who showed that high-wage routine workers used to migrate to abstract jobs, whereas low-wage routine workers to manual jobs. While Cortes and Groes et al. focused on the effect of the sorted migration on the composition of the labor force, we argue that this mechanism affects not only workers' *observable* characteristics, but also their *unobservable* characteristics, i.e., the distribution of skills within workers' groups, which appears to be the key element to understand the wage dynamics among low-earners.

The remainder of the paper is organized as follows. In Section , we present the May/ORG Census Samples database, and in Section 2.2 and 2.3, we provide some stylized facts regarding the relationship between occupational tasks and the evolution of wage inequality in the U.S. during the 1990s. In Section 2.4, we present the results of the CQR analysis, and in Section 3 we show how to rationalize the empirical findings using a simple model of the labor market endowed with RBTC operating though a labor-supply channel. Section 4 concludes.

2 Empirical Evidence

2.1 Data

To analyze the U.S. wage distribution, we employ the database constructed by Autor, Katz and Kearny (2008) [AKK] that combines two data sources commonly used in this literature. The first is the annual collection of the March Issues of the Current Population Survey (CPS), supplemented by data from the May Issues and the Outgoing Rotation Group, which provides a measure of weekly wages for the entire distribution of worked hours included in CPS surveys for the years from 1986 to 2002. We refer to this source as May/ORG CPS.[8] The second is the Fourth Edition (1977) and the Revised Fourth Edition (1991) of the U.S. Department of Labor's Dictionary of Occupational Titles [DOT]. Data from the May/ORG CPS are merged with the DOT to build a map between the occupations listed in the May/ORG CPS and their contents in terms of primary comparable *tasks*. The resulting database provides a panel of observations at the worker level with information regarding worker's occupation, his/her weekly wage the and corresponding wage percentile, the tasks performed on duty, and several socioeconomic characteristics. In the empirical analysis, we do not pursue a task-based classification of workers because no occupation implies performing one single task; therefore, there is not a unique correspondence between workers and task-types (e.g. "routine" or "manual" workers).[9] To avoid arbitrary assumptions, we perform the empirical analysis by using the distribution of task intensities across wage percentiles.

Following ALM (2003), we aggregate the original 44 tasks defined in the DOT into the following five groups of tasks: (i) EYEHAND, which is the ability to move hands and feet in coordination with the other senses, notably sight. These tasks are usually defined as *manual* in the literature; (ii) FINGDEX, which is finger dexterity. This group evaluates the ability to do something manual with skill and speed and consists of what is typically defined as *routine* tasks; (iii) STS, which is the ability to set limits, tolerances, or standards for any production process and consists mainly of routine tasks; (iv) DCP, which is the ability to undertake direction, control and planning – and involves the attitude to accept responsibility – for supervising and planning activities. Tasks in this group are typically

[8] Autor, Katz and Kearny (2008) and Lemieux (2006) provide a full set of descriptive statistics on these data. In particular, Lemieux pointed out the following drawbacks of the May/ORG CPS: (i) the treatment of censored wages, particularly top-coded wages; (ii) the existence of allocated or imputed wages for workers who do not respond to the survey; and (iii) the comparison of wages pre and post 1994, when several changes were implemented in the design of the survey. Autor, Katz and Kearny showed that the inclusion of data from the CPS March Issues help to address some of these issues. In this paper, we follow their empirical strategy in treating the data and do not address the remaining issues.

[9] For instance, although a clerk performs primarily routine tasks, such as making copies or performing calculations, he also performs manual tasks (answering phone calls) and cognitive tasks (taking minutes of meetings). Executives perform mostly cognitive tasks, such as organizing the firm's business, but they are also involved in routine tasks, such as checking variations in sales data.

defined as *cognitive* in nature; (v) MATH, which refers to general education, analytical and mathematical skills and the ability to engage in problem solving. It identifies the most typical *cognitive* tasks.[10] We adopt the five-group classification originally used in Autor, Levy and Murnane (2003) instead of the three-group classification (manual, routine, cognitive) used in more recent literature, (e.g. Autor and Dorn, 2013), because in side estimations we find that the five-group classification is more effective than the three-group classification in identifying the dynamics of the remunerations for the different types of tasks. In particular, we noted that STS behaves very differently from FINGDEX, and using a routine meta-group that includes both would add noise to the data and obscure the results. The same occurs with DCP and MATH, whose remunerations appear to have both different evolutions over time and different effects on wage inequality.

2.2 Price and Composition effects with tasks

As noted by Lemieux (2006), changes in wage inequality can be caused by (i) changes in the remuneration for workers' characteristics and/or (ii) changes in the distribution of workers' characteristics. In particular, growing wage inequality can be determined not only by higher skill premia, but also by increases in the employment share of workers with characteristics that intrinsically imply higher wage dispersion. For properly assessing the effect of skills remunerations on wages, Lemieux suggested a reweighing kernel approach to simulate counterfactual changes in the wage distribution while maintaining the labor force composition constant and thereby isolating the so-called *price effect* (effect of changing wages) from the *composition effect* (effect of changes in the labor force composition). Table 1 presents the results of the reweighing kernel analysis as applied to our data. The first panel reports overall, residual, and composition-adjusted wage variance for the initial and final sample periods, together with their changes in absolute terms.[11] The second panel shows the same statistics computed separately for above and below the median wage. Resid-

[10]To control for possible changes in the content of the tasks of each occupation across different periods of time, the original measures of tasks provided by the DOT is transformed into percentile values in ranking the task distribution in the initial year of the DOT (1960). As argued by ALM, 1960 can be safely assumed as the benchmark year because it was a year before the beginning of the implementation of computer practice in business and production.

[11]For the sake of consistency with the empirical analyses in following sections, we compute the statistics for the initial period by pooling the first three years of the sample (1986–1988), and that of the final period using the last three years (2000–2002).

ual wage variance is used as measure of RWI and is obtained from standard Mincer-type wage regressions repeatedly estimated for each year of the sample in which the log weekly wage is regressed on education, age and their cross-products, plus the intensity of each occupational-tasks group as defined in the previous section. Composition-adjusted wage variance is obtained from residual wage variance by combining the actual price function with the composition function from the initial year.

Table 1: Wage variance for male workers

	Overall wage variance	Residual wage variance	Composition-adjusted wage variance
1986/1988	0.290	0.175	0.175
2000/2002	0.315	0.180	0.194
Change	0.0245	0.0049	0.0181
percentiles	Change above and below the median		
below 50th	−0.0164	−0.0098	−0.0102
above 50th	0.0201	0.0228	0.0115

Overall wage variance is computed as the weighted variance among individual wages. *Residual wage variance* is computed as the weighted variance among residual wages, in which residuals come from the regression of log hourly wages over workers' education, age and the intensity of occupational task groups as defined in Section . *Composition-adjusted wage variance* is computed with the reweighing kernel approach described in Lemieux (2006) using the initial period as base year.

As Table 1 illustrates, controlling for the composition effect does not reduce the observed increase in RWI. As a matter of fact, the composition-adjusted wage variance increases more than the residual wage variance, suggesting that the price effect is the key factor explaining the observed change of RWI. The difference between our results and those of Lemieux (2006) is possibly due to the different definition of homogeneous workers used by Lemieux, which did not include occupational tasks (or workers' occupations) in wage regressions. Given that RBTC generates salary premia in some jobs and losses in others, it is reasonable to posit that the effects of tasks remunerations average out when workers with heterogenous occupations are included in the same groups. By inspecting the last two lines of Table 1, we also learn that the price effect has a key role in explaining wage inequality in the left hand side of the wage distribution, whereas price and composition effects are equally important in shaping the right hand side of the wage distribution.

2.3 The influence of tasks on RWI

In this section, we seek to understand whether occupational tasks can be suitable candidates for explaining RWI. To this end, we regress the total growth of RWI over the considered sample on task intensities. As in previous section, RWI is computed for narrowly defined groups of workers using the following characteristics: education, experience, gender and tasks groups. Because some of the 6,700 cells of homogeneous workers defined by these characteristics have only a limited number of observations, we build a pseudo-panel by pooling 1986–1988 as the initial period and 2000–2002 as the final period. Accordingly, in the specification (1) V_i is measured as the growth rate of residual wage variance in each cell between the initial and the final periods.

By denoting ehf_i, fgx_i, sts_i, dcp_i, and $math_i$ as the intensity of EYEHAND, FINGDEX, STS, DCP, and MATH in cell i, respectively, we estimate the following empirical specification:

$$V_i = \alpha + \beta_1 ehf_i + \beta_2 fgx_i + \beta_3 sts_i + \beta_4 dcp_i + \beta_5 math_i + \varepsilon_i \qquad (1)$$

where task remunerations are given by coefficients $\beta_j \ for \ j \in (1, ..., 5)$. We repeat the estimation four times. The first estimation is performed using the entire wage distribution, the second using the left tail only (wages below the 30th percentile), the third uses the middle echelon (from the 30th to the 60th percentile), and the fourth uses the right tail only (above the 60th percentile). Each estimation is repeated twice, either including or not including a number of additional variables to control for the effect of other factors that might affect wage inequality, namely, membership in unions, marital status and race (see Firpo, Fortin and Lemieux, 2011).

Table 2: The impact of occupational tasks on RWI growth

	Percentiles							
	all		below 30th		30th-60th		above 60th	
nonroutine manual	0.007	0.006	0.016**	0.016**	0.004	0.004	-0.006	-0.006
	(0.007)	(0.007)	(0.008)	(0.008)	(0.007)	(0.007)	(0.009)	(0.009)
routine manual	-0.018*	-0.019*	-0.021*	-0.021*	-0.003	-0.004	-0.021**	-0.021**
	(0.010)	(0.010)	(0.013)	(0.013)	(0.012)	(0.012)	(0.010)	(0.010)
routine cognitive	0.002	0.001	-0.011***	-0.011***	-0.004	-0.005	0.014***	0.014***
	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
nonroutine interactive	-0.003	-0.003	-0.011*	-0.010*	-0.003	-0.003	0.002	0.002
	(0.004)	(0.004)	(0.006)	(0.006)	(0.004)	(0.004)	(0.004)	(0.004)
nonroutine analytic	0.062***	0.065***	0.078***	0.078***	0.060***	0.062***	0.031***	0.032***
	(0.007)	(0.007)	(0.009)	(0.009)	(0.008)	(0.008)	(0.007)	(0.008)
dunionmme		-0.218***		-0.203**		-0.116**		-0.061
		(0.067)		(0.089)		(0.053)		(0.054)
dnonwhite		0.028		0.032		-0.027		-0.019
		(0.078)		(0.053)		(0.057)		(0.072)
dmarried		0.043		-0.043		-0.017		-0.006
		(0.061)		(0.045)		(0.043)		(0.054)
Constant	-0.324***	-0.339***	-0.356***	-0.362***	-0.351***	-0.364***	-0.184***	-0.189***
	(0.036)	(0.037)	(0.041)	(0.041)	(0.041)	(0.042)	(0.039)	(0.040)
N. of groups	6815	6815	5557	5557	6307	6307	5785	5785

The estimation is performed in two steps. In the first step (wage regressions), the overall growth rate of wages (first difference of log wages between initial and final period) is regressed on: education, age dummies – used as a proxy for workers' experience –, occupations, and their cross products. In the second step, the sum V_i of squared residuals from wage regressions is calculated for each cell of homogeneous workers (same education, age *and* occupation), and is regressed on task intensities. Because the number of observations differs noticeably among cells, each cell is weighted in the estimation using the sum of the individual weights of the workers belonging to that cell, as assigned in the MAY/ORG CPS database. This estimation strategy ensures that all available information is efficiently used but that no observation is over-weighted with respect to its original survey weight.

The results reported in Table 2 yield several insights. We shall explore them in turn. First, the estimations performed using the entire distribution reveal that (i) non-routine analytic tasks have the only positive and significant coefficient among tasks groups, (ii) routine manual has the only negative and significant coefficient and (iii) all of the remaining coefficients are nonsignificant. Second, when estimations are performed using only the lower percentiles, the coefficient of non-routine manual tasks turns significant and positive, whereas that of routine cognitive turns significant and negative. Thus, both groups of technology-substitute routine tasks appear to have negative effects on wage inequality among low-earners, whereas technology-neutral manual tasks appear to have a mild but positive effect.[12] Finally, in the estimations performed using the middle percentiles no coefficient is significant except for non-routine analytic tasks, which confirms that the bulk of the wage dynamics occurs at the periphery of the wage distribution. In general, previous results seem to indicate that technology-complementary non-routine analytic tasks push wage inequality upward, whereas technology-substitute routine manual tasks push wage inequality downward. More importantly, the analysis shows that the effects of tasks vary in intensity and sign across the wage distribution. Therefore, in next section we properly assess these effects by means of a quantile regressions analysis.

2.4 Counterfactual Analysis

In the following, changes in RWI are measured as variations in wage gaps between selected percentiles of the wage distribution that occurred between the initial period (1986–1988) and the final period (2000–2002). We employ the same data used to perform the kernel reweighing analysis and the OLS regressions, but percentiles below the $5th$ and above the $95th$ are trimmed to wash out the noise that is typical in data at the extremes of the wage distribution. In the quantile regressions, the following covariates are used: education, experience (proxied by age), task intensities by groups as defined in Section , union membership, marital-status and race.[13] Tables 3-4-5 present the estimation outcome. For each percentile interval, the reported price effect (composition effect) represents the counterfac-

[12]This finding incidentally supports Autor and Dorn (2013), who argued that an increase in the demand for, and the relative wage of, non-routine manual tasks embedded in service occupations occurred at the expense of routine tasks.

[13]The estimation is performed using only male workers. Estimation results for females are available upon request. In general, none of the results presented in this section are undermined when all workers are included in the analysis.

tual change between percentiles that would have occurred if the quantities (coefficients) of all the covariates had remained fixed at their initial values, but their coefficients (quantities) had taken the final period values. In Tables 4 and 5, the reported contributions of single variables measure the counterfactual change between wage percentiles that would have occurred if only the coefficient of the variable at issue had varied and all the other coefficients and all the quantities had remained fixed at their initial values. In other words, this method isolates the effect on changes in wage inequality between the Xth and Yth percentiles that would have occurred if only the remuneration for worker's characteristics Z had changed.

Table 3: Counterfactual Decomposition of changes in wage gaps by quantiles

Aggregate Decomposition	Percentiles			
	5th-30th	30th-60th	60th-95th	5th-95th
Price effect	-4.20	-0.13	7.89	3.56
	(.758)	(.556)	(.84)	(1.081)
between-group	0.53	0.55	3.68	4.75
	(.72)	(.521)	(.73)	(1.192)
within-group	-4.73	-0.68	4.21	-1.19
	(.787)	(.319)	(.722)	(1.113)
Composition effect	1.61	1.40	0.76	3.77
	(1.35)	(.955)	(1.495)	(1.717)
Total	-2.85	1.06	8.74	6.95
	(1.4)	(.947)	(1.543)	(1.902)

Sample: 1986/89-2000/02. Standard Deviations (in parenthesis) obtained using bootstrap procedure with 200 draws.

From an aggregate perspective, the estimation results reported in Table 3 mirror the evidence from the previous literature. The overall $5th-95th$ wage gap rose by approximately seven percentage points (Acemoglu, 2002) and the increase was fully driven by a larger dispersion among high wages, which more than compensated for the reduction among low wages (Buchinsky, 1998). When considering the entire distribution, the positive composition effect explains a prominent fraction of the overall change (Lemieux, 2006). However, once the distribution is analyzed separately above and below the median, the change in RWI net of the composition effect is again positive in the upper echelon and negative in the bottom echelon of the wage distribution (AKK, 2008). As a matter of fact, the first line of Table 3 qualify AKK findings by showing that the bulk of wage inequality dynamics is concentrated below the $30th$ and above the $60th$ percentiles, whereas wage inequality among the middle percentiles appears relatively stable. Finally, the last column of Table 3 reveals that the largest contributor to the total increase in RWI is the between-group price effect, i.e., increased wage differentials among workers performing different occupational tasks. The CQR analysis also clarifies two important caveats that apply to this result. First, when the estimation is performed using the entire distribution and without distinguishing the between-group from the within-group price effect, then the composition effect becomes the largest contributor to the growing inequality because the positive between-group price effect is partially offset by the negative within-group price effect. This finding reconciles our results with those of Lemieux (2006). Second, the importance of the between-group price effect is fully determined by the behavior of high wages. When the estimation is performed focusing only on the lower percentiles, RWI is entirely due to a negative within-group price effect, which is only partially compensated for by a positive composition effect (first column), whereas the between-group price effect is rather small and nonsignificant. This finding suggests that the dispersion of wages among workers performing the same tasks is the key determinant to explain wage inequality among low-earners. On top of previous results, Table 3 also shows that the total variation of wage inequality is nonsignificant around the median wage (30th-60th wage gap). The estimation reveals that this result is due to the opposite signs of a negative within-group price effect and a positive composition effect. In general, the magnitude of variations in the middle percentiles are small compared with those at the periphery of the distribution. Finally, the price effect (strong) and the composition effect (mild) in the upper percentiles are both shown to be positive and to jointly explain the large increase in the 60th-95th wage gap.

Turning to the contributions of single variables, Table 4 shows that the effect of tasks on within-group wage inequality outweighs the effects of all the other variables both when analyzing the distribution as a whole or by echelons, and the only exception is the effect of education on high wages. The largest contributors to the within-group price effect are *routine manual* and *non-routine analytic* tasks groups with symmetric and opposite effects on wage inequality. In particular, the first reduced wage inequality located in the bottom percentiles, whereas the second increased wage inequality located in the upper percentiles. More in detail, the variation in the remuneration of routine manual tasks

Table 4: Counterfactual Decomposition of changes in wage gaps

Price effect within-group		Percentiles			
		5th-30th	30th-60th	60th-95th	Total
education		-0.37	0.05	2.11	1.80
		(.443)	(.231)	(.607)	(.803)
experience		-0.30	-0.15	-0.27	-0.72
		(.491)	(.272)	(.528)	(.922)
tasks		-5.65	-2.05	2.17	-5.53
		(2.006)	(1.126)	(1.922)	(3.901)
	nonroutine manual	-0.12	-0.31	-0.73	-1.16
		(.301)	(.192)	(.399)	(.587)
	routine manual	-7.19	-3.55	-1.57	-12.31
		(2.141)	(1.166)	(1.977)	(4.126)
	routine cognitive	-1.26	-0.15	-0.33	-1.75
		(.591)	(.271)	(.494)	(.957)
	nonroutine interactive	-0.20	-0.18	0.59	0.21
		(.337)	(.18)	(.652)	(.759)
	nonroutine analytic	2.10	2.14	4.18	8.42
		(.563)	(.424)	(1.049)	(1.414)
union		-0.26	0.28	0.56	0.59
		(.129)	(.12)	(.25)	(.263)
married		-0.03	0.41	0.51	0.88
		(.286)	(.185)	(.378)	(.513)
race		0.28	0.06	-0.03	0.31
		(.19)	(.079)	(.123)	(.228)

Sample: 1986/89-2000/02. Standard Deviations (in parenthesis) obtained using bootstrap procedure with 200 draws.

implied a reduction in the 10th-30th wage gap of -7.2%, and this effect alone basically drove the entire evolution of wage inequality in the bottom echelon of the wage distribution. The variation in the remuneration of the non-routine analytic tasks group implied instead an increase of 4.2% in the 60th-95th wage gap, which accounts for half of the inequality variation in the upper echelon of the wage distribution. Finally, Table 5 reveals that the group of non-routine analytic tasks shares the strongest positive impact on between-group wage dispersion among high earners with education, which is consistent with the predictions of standard theories of human capital accumulation and the empirical findings of Piketty and Saez (2003). Regarding the other groups of tasks, the effects of non-routine manual, non-routine interactive and routine cognitive tasks on within-group wage inequality are barely significant; therefore, the overall effect of tasks on within-group wage inequality (-5.5%) is basically given by the difference between the negative effect of routine manual tasks (-12%) and the positive effect of non-routine analytic tasks ($+8\%$).

Table 5: Counterfactual Decomposition of changes in wage gaps

Price effect between-group		Percentiles			
		5th-30th	30th-60th	60th-95th	Total
education		0.66	0.85	2.99	4.49
		(.321)	(.299)	(.585)	(.727)
experience		-1.64	-1.12	-1.08	-3.84
		(.488)	(.34)	(.401)	(.606)
tasks		1.41	0.97	2.02	4.40
		(.557)	(.376)	(.646)	(1.005)
	nonroutine manual	-0.16	-0.21	-0.38	-0.75
		(.286)	(.187)	(.234)	(.324)
	routine manual	0.04	0.05	0.28	0.37
		(.249)	(.172)	(.303)	(.358)
	routine cognitive	0.25	0.14	0.21	0.60
		(.367)	(.243)	(.389)	(.457)
	nonroutine interactive	-0.79	-0.67	-1.03	-2.48
		(.268)	(.293)	(.44)	(.634)
	nonroutine analytic	2.12	1.61	2.88	6.61
		(.48)	(.355)	(.514)	(.797)
union		-0.11	-0.08	-0.02	-0.21
		(.104)	(.106)	(.164)	(.176)
married		0.12	0.08	0.09	0.28
		(.159)	(.11)	(.147)	(.257)
race		-0.10	-0.08	-0.09	-0.26
		(.205)	(.141)	(.183)	(.235)

Sample: 1986/89-2000/02. Standard Deviations (in parenthesis) obtained using bootstrap procedure with 200 draws.

Broadly speaking, previous results appear in line with the predictions of the routinization hypothesis. Changes in remunerations for occupational tasks are overwhelming in determining wage inequality and, in general, wage differentials noticeably grew along the wage distribution, eventually representing the single most important source of changes in wage inequality during the considered period. However, from the analysis performed separately on different echelons of the wage distribution, we also learn that this result is fully driven by the effect of wage differentials among high-earners, whereas the evolution of wage inequality among low-earners is almost fully determined by the contraction of wage dispersion *within* groups of homogeneous workers and, in particular, workers performing routine tasks. The fact that occupational tasks play almost no role in changing wage differentials among low-earners appears to conflict with the notion of RBTC operating through a labor-demand channel. For this reason, in next section we investigate an alternative formulation of the theory.

3 The Model

We consider a partial equilibrium model of the labor market in which a continuum of uniformly distributed income-maximizing individuals indexed $i \in [0, 1]$ are each endowed with an idiosyncratic level of ability a_i. Each individual inelastically supplies one unit of time to the labor market and decides which occupation to undertake. Following the job polarization literature, we assume that there are three occupations that differ in their degree of complementarity with technology: abstract jobs, h_t, routine jobs, z_t, and manual jobs, l_t.

3.1 Labor demand and wages

The demand for labor is formed by a price-taker firm that combines worked hours from each occupation into units of productive labor. Building on Galor and Moav (2000), we define the complementarity between occupations and technology in terms of an *erosion effect*. When technical innovations occur, they erode the number of jobs that are *not* complementary to technology either because innovations replace workers with machinery or – so long as occupations are substitutes – because innovations reduce the relative efficiencies of technology-neutral and technology-substitute labor compared with technology-complementary labor. The erosion effect is assumed to depend on the growth rate of technology, implying that the relative demand of occupations only changes when there are new waves of technical innovations, whereas it is constant along a stably increasing path of technology.[14] Moreover, we distinguish among occupations by assuming that workers whose occupations require a minimum level of ability in the model are more productive than other workers. The following composite labor aggregate H_t conveniently accommodates the previous assumptions:

$$H_t = \beta h_t + \beta(1 - \delta g_t)z_t + (1 - \delta g_t)l_t \tag{2}$$

where $g_t = (A_t - A_{t-1})/A_{t-1}$ is the growth rate of technology and A_t its level, $\delta \in (0, 1)$ measures the intensity of the erosion effect and $\beta \in (1, \infty)$ captures the extra productivity of skilled labor compared with unskilled labor. In the model, only abstract occupations are assumed to be complementary to technology; accordingly, both routine z_t and manual l_t occupations are subject to the erosion effect. In addition, the same erosion intensity is assumed for routine and manual occupations independently of their degree of complementarity/substitutability with technology. This strategy imposes less structure on the model and avoids arbitrary assumptions required to calibrate different values of δ, which may direct the results. The demand of routine labor is differentiated from that of manual labor because routine workers are rewarded with the skill premia β.

The cost-minimizing firm produces under a standard Cobb-Douglas technology and raises capital and hires labor in perfectly competitive markets. When the interest rate is constant, this formulation implies that the optimal ratio of capital to labor k_t is also constant, and the wage rate can be expressed as $w_t = A_t\bar{w}$, where $\bar{w} = f(\bar{k}) - f'(\bar{k})\bar{k}$, $f(\bar{k})$ is the production function, and f' is its first derivative.[15] Accordingly, the occupation-

[14]Galor and Moav showed that a formulation of the erosion effect in terms of growth rates – as opposed to the level of technology – helps disentangle the short-run from the long-run effects of technical change. In our partial-equilibrium setup, this formulation is mostly convenient for accommodating different effects of technical change on the labor demand (affected by the level and the growth rate of technology) and the supply of labor (affected only by the growth rate of technology).

[15]To analyze the general equilibrium implications of technical change, the labor market model presented in this section can be readily extended to a dynamic general equilibrium framework by considering a small open economy in which the representative firm sells its product for investment and consumption purposes to a continuum of lifetime utility-maximizing households. Under the standard assumptions of concavity, non-satiability and separability of the utility function – in addition to the assumption that different types of labor entail the same levels of disutility – the household's problem of occupational choice is separable from saving/consumption choices, and therefore coincides with the choice analyzed here.

specific wages are:

$$\begin{aligned} w_t^h &= \beta A_t \bar{w} & (3) \\ w_t^z &= \beta(1-\delta g_t) A_t \bar{w} & (4) \\ w_t^l &= (1-\delta g_t) A_t \bar{w} & (5) \end{aligned}$$

The wage rates (3)−(5) conform to the empirical evidence presented in previous section, in which abstract jobs represent the model counterparts of cognitive tasks. Technology is assumed to push abstract wages upward, whereas its effect on routine and manual wages is undetermined although surely smaller than the one on abstract wages.[16] As a consequence, technical change always enhances wage differentials between abstract and routine occupations, $[w_t^h/w_t^z = (1-\delta g_t)^{-1}]$, and between abstract and manual occupations $[w_t^h/w_t^l = \beta(1-\delta g_t)^{-1}]$, but leaves the wage differential between routine and manual occupations unaffected. This feature of the model shuts down the labor-demand channel in the bottom half of the income distribution, thus leaving the effect on the labor supply as the only channel through which technical change affects low wages, in line with the central concern of our analysis.

3.2 Labor Supply and Occupational Choices in equilibrium

The amount of efficiency units of labor that each individual can supply in each occupation depends on her ability and on the technological environment. In particular, technological progress is assumed to erode existing job skills, which can be reestablished by individuals through a learning process. This assumption is introduced in the model using

[16]The marginal returns of technical change to wage rates are:

$$\begin{aligned} \partial w_t^h/\partial g_t &= \beta \bar{w} A_{t-1} > 0 \\ \partial w_t^z/\partial g_t &= \beta \bar{w}\,(A_{t-1}(1-\delta g_t) - \delta A_t) \gtrless 0 \\ \partial w_t^l/\partial g_t &= \bar{w}\,(A_{t-1}(1-\delta g_t) - \delta A_t) \gtrless 0 \end{aligned}$$

The effect of g_t on w_t^z and w_t^l is undetermined because technical change generates two countervailing forces. The higher level of technology A_t raises the marginal productivity of all types of labor, whereas the higher growth rate g_t erodes the occupation-specific demands of labor, thus pushing the equilibrium wage rate downward. The overall effect will depend on which of these two forces prevails in equilibrium.

the linear formulation suggested by Galor and Moav,

$$h_t^i = a_i - (1 - a_i)g_t \quad (6)$$

$$z_t^i = 1 - (1 - a_i)g_t \quad (7)$$

$$l_t^i = 1 \quad (8)$$

Equations (6) and (7) posit that the number of the efficiency units of abstract and routine labor increases with a_i, which replicates the assumption that ability reduces the cost of learning [Bartel and Sicherman, 1998] and decreases with $g_t > 0$ because we assume that workers must devote a fraction of their time to learning innovations to maintain their supply of skilled labor at a constant level. Equation (8) posits that manual occupations require no learning processes, which replicates the standard assumption that working duties in manual occupations are technology-neutral [Autor, Levy and Murnane, 2003]. Therefore, the supply of manual labor is constant even in a changing technological environment and always coincides with workers' time endowment. Three features of the adopted formulation are worth emphasizing. First, learning costs in abstract and manual occupations $[(1 - a_i)g_t]$ are equal, which is a conservative assumption with respect to our results. In fact, any labor supply function entailing a learning advantage for abstract compared with routine occupations would strengthen the migration of routine workers to abstract jobs after innovations, which would enhance the impact of technology on the wage dynamics. Second, only abstract occupations reward ability in a stationary technological environment, which in our framework is a natural consequence of the assumption that only abstract occupations are complementary to technology in the model. Third, technology is less costly for manual than for routine workers, which is a consequence of the different relationships of RBTC with manual (neutral) and routine (substitute) occupations postulated by the routinization hypothesis [Autor, Levy and Murnane, 2003; Acemoglu and Autor, 2011].

Each individual chooses which occupation to undertake in seeking to maximize her income and by observing the wage rates (3)−(5) and learning options (6)−(8). Because the different types of labor are perfect substitutes (equation 2), individual i will choose the highest among the following earning possibilities:

$$I_{i,t}^{h} = w_t^h \cdot h_t^i = \beta w_t \left(a_i - (1-a_i)g_t\right) \tag{9}$$

$$I_{i,t}^{z} = w_t^z \cdot z_t^i = \beta(1-\delta g_t) w_t \left(1-(1-a_i)g_t\right) \tag{10}$$

$$I_{i,t}^{l} = w_t^l \cdot l_t^i = (1-\delta g_t) w_t \tag{11}$$

As equations (9)−(11) illustrate, the marginal returns of ability to income are highest for abstract and lowest for manual occupations, $[\partial I_t^h/\partial a_i = \beta w_t(1+g_t) > \partial I_t^l/\partial a_i = \beta w_t g_t(1-\delta g_t) > \partial I_t^z/\partial a_i = 0]$, therefore determining in equilibrium a sorted mapping between occupations and income in which the most capable individuals obtain abstract jobs and land in the top echelon of the income distribution, less capable individuals obtain routine jobs and are situated in the middle echelon and the least capable individuals obtain manual jobs and are found in the lowest echelon. By equating pairwise the earning options for individual i, we can characterize the parametric values for the thresholds of occupational-switching ability as follows

$$a^* = 1 - \frac{1-\beta^{-1}}{g_t} \tag{12}$$

$$a^{**} = \frac{1-\delta g_t + \delta g_t^2}{1+\delta g_t^2} \tag{13}$$

Under certain conditions that guarantee the existence of all three types of labor in equilibrium,[17] the model argues that every individual with a level of ability above a^{**} will choose abstract occupations, those with a level of ability below a^* will choose manual occupations, and everyone in the middle will choose routine occupations. Figure 1 depicts the income and the associated occupational distribution in equilibrium as a function of individual ability. The upward frontier of income possibilities represents the overall supply of labor in equilibrium after occupational choices are made.

[17]Using equations (12) and (13), it can be shown that there *always exists* a calibration of $\{\beta\,,\delta\}$ that guarantees a positive mass of individuals in each occupation, i.e., $0 < a^* \leq a^{**} < 1$, and that this calibration must fulfill two conditions: $\beta(1-g_t) > 1$ and $\delta < \frac{\beta-1}{g_t^2}$. Intuitively, for routine occupations to exist, (i) routine income should be greater than manual income at least for the lowest level of ability, and (ii) the erosion effect should be not *too* large; otherwise, in presence of technological progress, every individual will find her supply of abstract labor high enough to choose abstract over routine occupations.

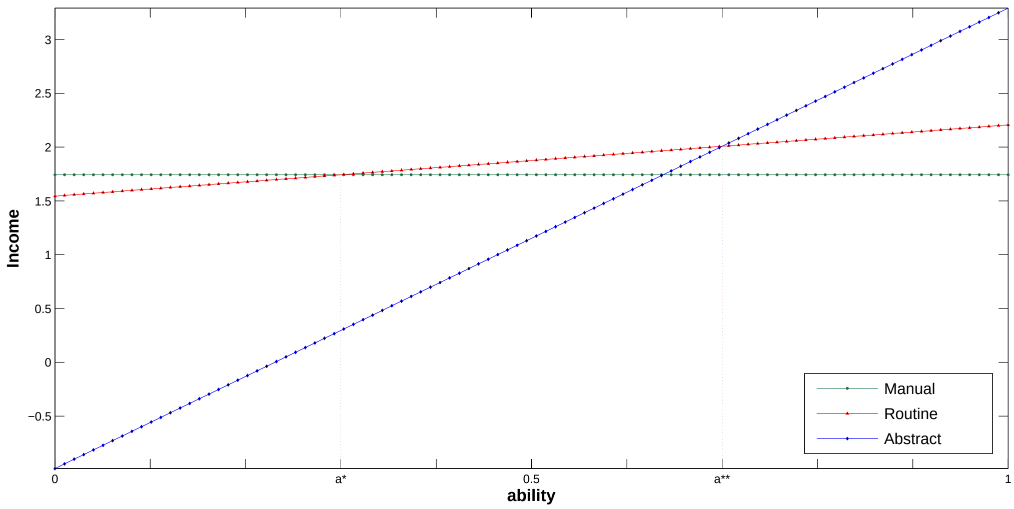

Figure 1: Income and Occupational Distributions

3.3 Technical change and the income distribution

The equilibrium characterized in previous section can be used to analyze the implications of technical change on the dynamics of income inequality. According to equations (9)−(11), a variation of g_t imply that both the labor demand (changing wages) and the labor supply (changing occupational choices) shift, thus determining a new equilibrium distribution of income and a new allocation of occupations along this distribution. In particular, we show that workers' inflows generated by technical change imply variations in the ability endowment for each occupation, thereby affecting within-group income inequality. These results are formally established in the next two propositions.

Proposition 1 (Employment) *Consider a partial equilibrium economy in which income-maximizing agents are endowed with heterogenous levels of innate ability. Assume that there are three types of labor whose relationships with individual ability are defined in equations (6)−(8). Each type of labor is hired by a profit-maximizing firm in a perfectly competitive labor market and employed in a constant returns to scale production function using the composite labor aggregate (2). In equilibrium,*

(i) whenever the growth rate of technology g_t increases, the mass of abstract workers univocally grows.

(ii) whenever the growth rate of technology g_t increases, the mass of routine workers univocally diminishes.

(iii) whenever the growth rate of technology g_t increases, the mass of manual workers univocally grows.

Proposition 1 replicates the well-known job polarization dynamics. The effect of technical change on the composition of the labor force is non-monotone along the income distribution, reducing employment in the middle echelon and increasing it at the extremes of the distribution. In our model, this mechanisms is generated by the twofold effect of g_t on the labor market. On the one hand, after an increase in g_t, wage differentials widen monotonically and reward abstract occupations, in particular. This condition pushes the most capable of routine workers to revise their occupational choices and to eventually switch to abstract jobs. The upward migration is limited to the fraction of routine workers possessing sufficiently high ability to guarantee a supply of abstract labor that makes the switch convenient. Otherwise, the routine worker maintains her current job. On the other hand, fast-flowing technical innovations generate a downward migration due to increased learning costs. If the ability level of a routine worker is low enough, when there is higher g_t, her supply of labor diminishes up to the point at which it becomes lower than the possible income from manual jobs. Eventually, all routine workers with an ability level below a certain threshold fail to catch up with the new technology and switch to manual occupations. It is worth noting that downward migration only occurs in one direction because higher levels of g_t have no effect on the relative wage between manual and routine jobs, although they reduce the supply of efficient units of routine labor. Thus, the expected income from routine occupations diminishes, the income from manual occupations remains constant and, as a result, no manual worker finds the switching option desirable.

One key feature of previous results is that workers' migration is *not* random but sorted along the ability distribution, involving routine workers placed at the extremes of the ability interval of routine occupations, either toward the bottom of the ability interval of abstract occupations, or toward the top of the ability interval of manual occupations. As a consequence, the ability intervals for each occupational group change in response to variations of technical change, eventually affecting the dispersion of wages within occupational groups. The following proposition establishes this result.

Proposition 2 (Ability distribution) *Consider the same economy defined in Proposition 1 and define the ability intervals of abstract, routine, and manual occupations, respectively, as:* $\overline{a}^h \equiv a_i \in (a^{**}, 1]$, $\overline{a}^z \equiv a_i \in [a^*, a^{**}]$, *and* $\overline{a}^l \equiv a_i \in [0, a^*)$. *Then, in equilibrium:*

(i) whenever the growth rate of technology g_t *increases satisfying* $g_t < \sqrt{\delta^{-1}}$, *then* $\overline{a}^h$ *widens, i.e., the dispersion of ability among abstract workers increases;*

(ii) whenever the growth rate of technology g_t *increases, then* $\overline{a}^z$ *univocally narrows, i.e., the dispersion of ability among routine workers in equilibrium diminishes;*

(iii) whenever the growth rate of technology g_t *widens, then* $\overline{a}^l$ *widens, i.e., the dispersion of ability among manual workers in equilibrium univocally increases.*

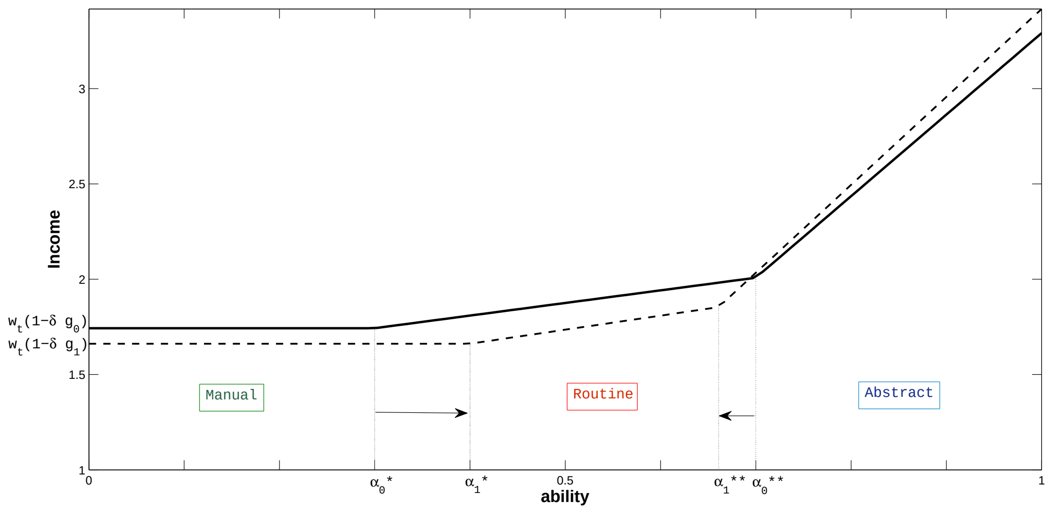

Figure 2: Within-group wage inequality as a function of technological progress.

The results in Proposition 2 have direct implications on within-group income inequality. When ability intervals increase (diminish), workers are less (more) homogeneous and the associated distribution of income within group becomes less (more) equal. Thus, Proposition 2 states that inequality increases with the pace of technological progress for abstract and manual workers, whilst it decreases for routine workers. This effect of technology on inequality *within* occupational groups can be directly observed in the model economy by defining the ratio of the highest to the lowest income in each group, i.e.

$$\sigma_t^h = \frac{I^h(a=1,g_t)}{I^h(a=a^{**},g_t)} \quad = \quad \frac{1+\delta g_t^2}{1-\delta g_t} \tag{14}$$

$$\sigma_t^z = \frac{I^z(a=a^{**},g_t)}{I^z(a=a^{*},g_t)} \quad = \quad \frac{\beta}{1+\delta g_t^2} \tag{15}$$

$$\sigma_t^l = \frac{I^l(a=a^{*},g_t)}{I^l(a=0,g_t)} \quad = \quad 1 \tag{16}$$

Equations (14)−(16) represent the model counterparts of the empirical within-group wage inequality analyzed in Section 2.4. Consistently with the provided empirical evidence, the effect of g_t is positive on σ_t^h and negative on σ_t^z. The effect on σ_t^l is null, but this result is a byproduct of the simplifying assumptions used in the linear model, which imply that income for manual workers is constant without regard to the worker's ability. It is worth noting that the effect of g_t on income inequality does not coincide with that on the ability intervals because the employed measure of inequality is not dimensionless; therefore, variations in the remuneration for ability has a direct effect on income inequality, although they have no direct effect on the ability intervals. In the cases of abstract and routine occupations, a variation in the price of ability operates on inequality in the same direction than the sorted migration, thereby univocally determining the overall impact of g_t. Instead, for manual workers the two effects operate in opposite directions, thus making the overall impact of g_t *a priori* undetermined. In general, within-group income inequality in the model closely mimics the one observed in the data and, in particular, variations in income inequality among low-earners are determined only by the response of the labor supply to RBTC, whereas the labor demand, i.e., the wage differential between manual and routine workers, remains fixed after technical innovations.

4 Conclusions

In this paper, we analyze the effect of RBTC on wage inequality among low, medium and high earners. Depending on which echelon of the wage distribution is considered, two different patterns unfold in the analysis. In the upper echelon, technical change appears to operate on wage inequality through both a labor-demand and labor-supply channel. By raising the remunerations of technology-complementary tasks with respect to the other occupational tasks, RBTC pushes upward (i) wage differentials among occupations comprising different occupational tasks, and (ii) the dispersion of wages within occupational groups, where the latter effect is typically intended to be generated by the increased remuneration for unobservable technology-complementary skills. In the bottom echelon, we find instead that RBTC has no effect on wage differentials, even though it continues to have a strong within-group price effect. We argue that this evidence points toward a model in which RBTC operates through a labor supply channel. To clarify this point, we develop a simple model of skill-heterogenous agents augmented with the routinization hypothesis in which agents face learning costs with respect to operating new technology. The model argues that technical change can have a non-monotone effect on wage inequality along the wage distribution similar to the one observed in the data even in absence of changes in wage differentials.

References

[1] Acemoglu, D. (2002): "Technical Change, Inequality, and the Labor Market", Journal of Economic Literature, 40(1), 7-72;

[2] Acemoglu, D. and Autor, D. (2011): "Skills, tasks and technologies: Implications for employment and earnings", Handbook of labor economics, 4, 1043-1171;

[3] Autor, D. and Dorn, D. (2013): "The growth of low-skill service jobs and the polarization of the US labor market", The American Economic Review, 103(5), 1553-1597;

[4] Autor, D. and Frank Levy and Richard J. Murnane (2003): "The Skill Content of Recent Technological Change: An Empirical Explo- ration", Quarterly Journal of Economics, 118(4), 1279-1333;

[5] Autor, D. and Katz, L. and Kearney, M. (2005): "Rising Wage Inequality: The Role of Composition and Prices", NBER Working Paper 11628;

[6] Autor, D. and Katz, L. and Kearney, M. (2008): "Trends in US wage inequality: Revising the revisionists", The Review of Economics and Statistics, 90(2), 300-323;

[7] Bartel, Ann P., and Nachum Sicherman (1998): "Technological Change and the Skill Acquisitions of Young Workers", Journal of Labor Economics, XVI, 718-755;

[8] Berman, E., Bound, J., and Griliches, Z. (1994): "Changes in the Demand for Skilled Labor within US Manufacturing: Evidence from the Annual Survey of Manufacturers", The Quarterly Journal of Economics, 109 (2), 367-397;

[9] Bound, J., and G. Johnson (1992): "Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations", American Economic Review, 82(3), 371-92;

[10] Buchinsky, M. (1998): "Recent advances in quantile regression models: a practical guideline for empirical research", Journal of human resources, 33 (1), 88-126;

[11] Cortes, G. (2016): "Where Have the Middle-Wage Workers Gone? A Study of Polarization using Panel Data", Forthcoming, Journal of Labor Economics;

[12] Di Nardo, J. and Fortin, N. and Lemieux, T. (1996): "Labor Market Institutions and the Distribution of Wages: a semiparametric approach", Econometrica, 64 (5), 1001-1044;

[13] Firpo, S. and Fortin, N. and Lemieux, T. (2011): "Occupational tasks and changes in the wage structure", IZA Discussion paper series, No. 5542;

[14] Freeman, R. B. and Katz, L. F. (1995): "Introduction and summary. In Differences and Changes in Wage Structures", 1-22. University of Chicago Press;

[15] Galor, O. and Moav, O. (2000): "Ability-biased technological transition, wage inequality, and economic growth", Quarterly Journal of Economics, 115 (2), 469-497;

[16] Goos, M. and Manning, A. (2007): "Lousy and lovely jobs: The rising polarization of work in Britain", The Review of Economics and Statistics, 89(1), 118-133;

[17] Goos, M. Manning, A., and Salomons, A. (2009): "Job polarization in Europe", American Economic Review papers and proceedings, 99, 58-63;

[18] Goos, M. Manning, A., and Salomons, A. (2011): "Explaining Job Polarization: The Roles of Technology, Offshoring and Institutions", Center for Economic Studies - Discussion papers, 11.34, Katholieke Universiteit Leuven;

[19] Groes, F. and Kircher, P. and Manovskii, I. (2015): "The U-Shapes of Occupational Mobility", Review of Economic Studies, 82(2), 659-692;

[20] Juhn, C. and Murphy, K. M. and Pierce, B. (1993): "Wage inequality and the rise in returns to skill", Journal of political Economy, 410-442;

[21] Katz, L. F. and Murphy, K. M. (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors", Quarterly Journal of Economics, 107(1), 35-78;

[22] Krueger, A. B. (1993): "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989", The Quarterly Journal of Economics, 108(1), 33-60;

[23] Lemieux, T. (2006): "Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill?", The American Economic Review, 96 (3), 461-498;

[24] Levy, F. and Murnane, R. J. (1992): "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations", Journal of Economic Literature, 30(3), 1333-81;

[25] Machado and Mata (2005): "Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression", Journal of Applied Econometrics, 20, 445-465;

[26] Mishel, L., Schmitt, J., and Shierholz, H. (2013): "Don't Blame the Robots: Assessing the job polarization explanation of growing wage inequality", EPI/CEPR Working Paper, Nov.19;

[27] Piketty, Thomas and Saez, Emmanuel (2003): "Income Inequality in the United States, 1913-1998", Quarterly Journal of Economics, 118(1), 1-39;

A The Counterfactual Quantile Regressions Approach

The CQR analysis and the associated wage inequality decomposition presented in Section 2.4 are obtained by performing the following steps. Let $Q_\theta(w_t|X_t)$ for $\theta \in (0,1)$ be the quantile θ^{th} at time t of the wage distribution conditional on a vector of k covariates x_t. First, we repeatedly estimate the Quantile Regressions

$$Q_{\theta_i}(w_t|X_t) = X_t'\beta_{\theta_i,t} \tag{17}$$

for $i = \{1,\ldots,10.000\}$ random draws of the quantile θ_i. For each draw, the procedure is repeated twice, for the initial and final period. Then, the estimated coefficients $\hat{\beta}_{\theta_i,t}$ are collected into a 10.000 rows vector $\hat{B}_t$ that is used to obtain the marginal density of w_t by generating a random sampling, $x^*_{i,t}$, from the rows of X_t and using

$$w^*_{i,t} \equiv x^*_{i,t}\hat{\beta}_{\theta_i,t} \tag{18}$$

as the unconditional distribution of wages. Finally, the simulated unconditional quantile, $\hat{\theta}$, is calculated as: $\hat{Q}_{\theta,t}(w^*_{i,t})$.

Next, the growth of wage inequality is defined as the change in the distance between two selected percentiles (θ and θ') and two selected periods (s and t), i.e. $(\hat{Q}_{\theta,s} - \hat{Q}_{\theta,t}) - (\hat{Q}_{\theta',s} - \hat{Q}_{\theta',t})$, or equivalently $\Delta\hat{Q}_{\theta,s,t} - \Delta\hat{Q}_{\theta',s,t}$. Now, given the median of the simulated distribution in year t, $\hat{Q}_{50,t}(\hat{\beta}_{50,T}, X_t)$, and defining $\hat{\beta}^{\omega}_{\theta,t} = \hat{\beta}_{\theta,t} - \hat{\beta}_{50,t}$ as the difference between the estimated coefficient in percentile θ and the median coefficient, for any percentile θ the change between two periods can be decomposed as follows:

$$\Delta\hat{Q}_{\theta,s,t} = \Delta\hat{Q}^{\omega}_{\theta,s,t} + \Delta\hat{Q}^{b}_{\theta,s,t} + \Delta\hat{Q}^{X}_{\theta,s,t} \tag{19}$$

where

$$\Delta\hat{Q}^{\omega}_{\theta,s,t} = \hat{Q}_{\theta}(\hat{\beta}_{50,s} + \hat{\beta}^{\omega}_{\theta,s}, X_s) - \hat{Q}_{\theta}(\hat{\beta}_{50,s} + \hat{\beta}^{\omega}_{\theta,t}, X_s) \tag{20}$$

is within-group wage change in percentile θ,

$$\Delta\hat{Q}^{b}_{\theta,s,t} = \hat{Q}_{\theta}(\hat{\beta}_{50,s} + \hat{\beta}^{\omega}_{\theta,t}, X_s) - \hat{Q}_{\theta}(\hat{\beta}_{50,t} + \hat{\beta}^{\omega}_{\theta,t}, X_s) \tag{21}$$

is the between-group wage change, and

$$\Delta\hat{Q}^{X}_{\theta,s} = \hat{Q}_{\theta}(\hat{\beta}_{50,t} + \hat{\beta}^{\omega}_{\theta,t}, X_s) - \hat{Q}_{\theta}(\hat{\beta}_{50,t} + \hat{\beta}^{\omega}_{\theta,t}, X_t) \tag{22}$$

is the composition effect. Eventually, the estimated overall wage inequality growth can be calculated using equation 19.

B Proofs

B.1 Proposition 1

Given the assumption that individuals are uniformly distributed on the [0, 1] continuous interval and holding the condition $0 \leq a^* < a^{**} \leq 1$, then item (i) directly follows from the fact that $\frac{\partial a^{**}}{\partial g_t} < 0$ and item (iii) from the fact that $\frac{\partial a^*}{\partial g_t} > 0$. Item (ii) follows from the contemporaneous increase in a^* and reduction in a^{**}.

B.2 Proposition 2

Using equations (12) and (13), the ability interval for each occupation can be written as function of technical change. That is:

$$\overline{a}^h = |1 - a^{**}| = \frac{\delta g_t}{(1 + \delta g_t^2)} \tag{23}$$

$$\overline{a}^z = |a^{**} - a^*| = \frac{1 - \beta^{-1}(1 + \delta g_t^2)}{g_t(1 + \delta g_t^2)} \tag{24}$$

$$\overline{a}^l = |a^* - 0| = \frac{\beta^{-1} - 1 + g_t}{g_t} \tag{25}$$

To prove the Proposition, it is sufficient to take the derivative of equations (23)-(25) w.r.t. g_t. Then, item (ii) and (iii) follow, respectively, from: $\frac{\partial \overline{a}^z}{\partial g_t} < 0$ and $\frac{\partial \overline{a}^l}{\partial g_t} > 0$. Finally, condition $g_t < \sqrt{\delta^{-1}}$ guarantees that $\frac{\partial \overline{a}^h}{\partial g_t} = \frac{\delta - \delta^2 g_t^2}{\left(1+\delta g_t^2\right)^2} > 0$, thus proving item (i).