Working papers series

WP ECON 15.01

Learning New Technology: the Polarization of the Wage Distribution

Manuel Hidalgo-Pérez (U. Pablo de Olavide)

Benedetto Molinari (U. Pablo de Olavide)

JEL Classification numbers: J24, J31, O33

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

Department of Economics

Learning New Technology: The Polarization of the Wage Distribution*

Manuel Hidalgo-Pérez
Universidad Pablo de Olavide

Benedetto Molinari
Universidad Pablo Olavide

December 18, 2014

Abstract

This paper presents novel evidence regarding the relationship between technological progress, occupational tasks and wage inequality. By applying a counterfactual quantile regression analysis to historic U.S. data, we show that the evolution of wage inequality in the lower echelon of the wage distribution was due entirely to a reduction of within-group wage inequality, which was determined, in turn, by more homogeneous remuneration paid to workers performing routine tasks. Changes in the differential between the remuneration paid to technology-complementary and technology-substitute tasks had only a negligible impact on wage inequality among low-wage workers, which casts some doubt on the validity of basing a theory of wage inequality on routinization-biased technical change operating through a labor demand channel. To reconcile the routinization hypothesis with the data, we develop a model in which skill-heterogeneous workers face endogenous occupational choices and learning costs in connection with operating a new technology. Even in the absence of changes in wage differentials, the model argues that technical change can generate an empirically consistent non-monotone effect on wage inequality by affecting the average level of skills within different groups of workers.

*We wish to thank and acknowledge useful comments regarding this paper from David Dorn, Antonio Villar, Nacho Garcia-Pérez, the participants at the 28th Annual Congress of the European Economic Association (EEA), the Spanish Economic Association Annual Meeting (SAEe), the Macroeconomic Workshop at CEMFI, the X Meeting of the Spanish Association of Labor Economics (AEET), and the XXIX AIEL National Conference. Corresponding author: Manuel Hidalgo-Peréz, *mhidper@upo.es*, Department of Economics, Universidad Pablo Olavide, Sevilla (Spain – EU).

JEL Classification: J24, J31, O33

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

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 [with respect to wage inequality, see Juhn, Murphy and Pierce (1993) for the 1980s, Acemoglu (2002) and Lemieux (2006) for the 1990s, and Acemoglu and Author (2010) for the 2000s.] This evidence stimulated a substantial debate that has continued in the literature 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 that was biased toward skilled employment [Juhn, Murphy and Pierce (1993)].[1] Although the latter became the mainstream explanation, several empirical studies have questioned its role as a determinant of growing 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 criticism, several authors have put forward a more nuanced version of SBTC to reappraise the relationship between technology and wage inequality. Building on the routinization hypothesis of Author, Levy and Murnane (2003), who studied the effects of firms' adoption of computer technology, ICT and automated machines on the labor market, the job polarization theory suggests that technology affects wage inequality by widening wage differentials among groups of workers who perform different tasks.[2] Because workers have multiple skills with which they perform

[1]The skill-biased technical change [SBTC] hypothesis is based on 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 skill-biased by Bound and Johnson (1992), Katz and Murphy (1992), Levy and Murnane (1992), Krueger (1993), Berman et al. (1994), among others.

[2]see Acemoglu and Author (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 Author (2011) provide an exhaustive analysis of the empirical facts that the canonical SBTC hypothesis cannot explain but that a task-based model can.

different tasks – and because technology is complementary only for some and acts as a substitute for others – technical change is expected to raise the wages of workers performing technology-complementary tasks, while reducing the wages of workers performing technology-substitute tasks.[3] Moreover, 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. Thus, technical change is expected to concurrently depress the remuneration of middle-earners and raise the remuneration of high-earners, which thus explains the increase in wage inequality through a polarization process of the wage distribution.[4]

Until now, the empirical literature has only provided circumstantial evidence regarding the relationship between job polarization and wage inequality [Mishel, Shierholz and Schmitt (2013)]. Autor, Katz and Kearny (2008) [AKK] were the first to address this issue using systematic data from the U.S. labor market and made two findings when they clustered workers according to their education, experience and gender. On the one hand, since the late 1980s, residual wage inequality [RWI][5] varied in opposite directions above and below the median wage, even when controlling for the composition effect [Lemieux, 2006]. In particular, RWI increased above and diminished below the median, thus mimicking the behavior of overall wage inequality that was first shown by Buchinsky (1998). On the other hand, AKK showed that (i) employment diminished in jobs mainly consisting of technology-substitute routine tasks, which are placed in the lower echelon of the skill distribution, and (ii) employment increased in jobs consisting mainly of technology-complementary cognitive tasks, which are placed in the upper echelon.[6] Hence, these authors concluded that the RWI dynamics are compatible 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.

By pursuing an approach based on RIF-Regressions, Firpo, Fortin and Lemieux (2011) (FFL) later provided a direct estimation of the overall effect of occupational tasks on wage inequality. Af-

[3]For instance, technological progress replaced (skilled) "blue collar" workers in manufacturing plants with automated machines, which enhanced the productivity of "white collar" officers in charge of the assembly lines. Technical change thus generated a comparative advantage of cognitive skills compared with manual dexterity that eventually implied an increase in the wages of white collar workers relative to blue collar workers.

[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.

[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.

ter controlling for institutional and socio-economic explanatory variables (including minimum wage, de-unionization, age and education), these authors claimed 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*" whereas offshorability became the key factor beginning in the 2000s. This conclusion, however, does not raise a unanimous consensus in the literature. The routinization hypothesis predicts that the remuneration of technology-complementary and technology-substitute tasks should vary in opposite directions. Therefore, if the theory is correct, the direct effects of technical change on wage inequality should be revealed in the behavior of wage differentials among workers performing different occupational tasks. However, in an analysis of the evolution of wages for 250+ detailed occupations, Mishel, Shierholz and Schmitt (2013) find that the effects of wage differentials on wage inequality is too mild to support the job polarization theory's explanation for growing wage inequality.

In this paper, we attempt to reconcile previous findings by seeking different channels through which changes in the remuneration of occupational tasks operate on wage inequality. We group workers according to the tasks performed on duty and then pursue a decomposition approach that separately identifies the price effect of occupational tasks on *between-group* and *within-group* wage inequality. The job polarization theory implies that the so-identified between-group price effect is a valid indicator of RBTC operating through a labor demand channel, and the signs of the effects of technology-complementary and technology-substitute tasks on between-group wage inequality can be used to test 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 (DOT), which is the same source of data used by AKK. This database is firstly used to provide suggestive evidence about the relationship between wage inequality and occupational tasks. By employing the reweighing kernel approach of Lemieux (2006), we identify the contributions of price and composition effects to RWI. When occupational tasks are maintained constant in the reweighing analysis, we show that the contribution of the composition effect appears smaller than has been found in the previous literature, and the price effect appears to be the crucial determinant of the observed evolution of RWI. Hence, we test whether occupational tasks are suitable candidates for explaining wage inequality by performing a set of simple OLS regressions in which the change in RWI over the sample is regressed on task intensities. The results indicate that occupational tasks significantly affect RWI, and the signs of the effects for different task groups conform with the theoretical predictions.

Based on previous evidence, we assess the effect of occupational tasks on between-group and within-group wage inequality by performing the extension proposed by Author, Katz and Kearny (2005) of Machado and Mata (2005) counterfactual quantile regressions [CQR]. Compared with other approaches used to analyze wage inequality in the literature, the CQR approach features two key properties exploited in our analysis. First, it disentangles the *price effect* from the *composition effect* (as in standard Oaxaca-Blinder procedures), and it also permits the separation of the price effect that impacts between-group and within-group wage inequality. Second, it identifies the contributions of single variables in determining the price effect on both a within-group and a between-group basis.[7] When applied to U.S. labor market data over the 1986-2002 period, the CQR analysis reveals that the decline of wage inequality in the lower echelon of the wage distribution (below the 30th percentile) is explained almost entirely 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 workers – same education, experience *and* tasks performed on duty – appear to have marginal effects on wage dispersion. On the contrary, the dynamics of wages in the upper echelon of the distribution (above the 60th percentile) appears to be explained in equal parts by positive composition and price effects. In addition, we find that both the within-group and between-group price effects are almost entirely 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. Whereas the estimation outcome for the upper echelon of the wage distribution is consistent with the predictions of the standard job polarization theory, the estimation outcome for the lower echelon appears in contradiction with it. If remuneration for the skills used to perform routine tasks has diminished, we should observe that routine tasks negatively affect between-group wage inequality due to lower wage differentials between workers performing routine tasks and the other workers, which in the lower echelon of the wage distribution are mainly manual workers whose initial wages are lower than those of routine workers. Instead, the results show that routine tasks have an almost negligible effect on between-group wage inequality. We argue that this evidence does not necessarily point toward the dismissal of the routinization hypothesis, but it may indicate

[7] Compared with the RIF-regression approach followed by FFL, the CQR is disadvantaged because it does not allow the identification of the overall effect of tasks on wage inequality. However, this issue is not the focus of our analysis because we are training our attention on understanding the channels through which technology operates on wages instead of focusing on providing an overall assessment of the effect of technical change.

a model in which RBTC affects the wage distribution through a labor supply channel. We elaborate upon this point 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. As is customary in these types of models, labor costs are assumed to be increasing at the pace – and decreasing in the ability level – of technical change. In addition, we maintain the wage differential between manual and routine occupations constant to shut down the labor demand channel. Under these conditions, the model demonstrates that technical change generates a reduction of within-group wage inequality in the lower echelon and an increase of skill distribution in the upper echelon that is akin to the evidence observed in the actual data. Such dynamics are generated by a sorted migration of 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 the data demonstrate. The existence of a *sorted* migration of routine workers to manual and abstract occupations has been empirically supported by Cortes (2012) and Manovskii (2014), who showed that high-wage routine workers migrated to abstract jobs, whereas low-wage routine workers migrated to manual jobs. Cortes and Manovskii focused on the impact of the sorted migration on wage inequality though the composition effect. In this paper, we argue that the same mechanism affects not only the composition of the labor force but also the distribution of income within occupational groups and that the latter appears to be the key element to understand low-wage dynamics.

The remainder of this paper is organized as follows. In Section 2.1, we present the May/ORG Census Samples database, and in Section 2.2 and 2.3, we elaborate upon some suggestive evidence regarding the relationship between occupational tasks and the evolution of wage inequality in the U.S. during the 1990s. In Section 2.4, we provide 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 and the routinization hypothesis. Section 4 concludes.

2 Empirical Evidence

2.1 Data

To analyze the U.S. wage distribution, we employ the database constructed by AKK that combines two sources of data commonly used in the literature on job polarization. The first data used are from the March Issues of the Current Population Survey (CPS), supplemented by data from the May Issues and the Outgoing Rotation Group, to provide 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 source of data is the Fourth Edition (1977) and the Revised Fourth Edition (1991) of the U.S. Department of Labor's Dictionary of Occupational Titles [henceforth, 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 consisting of data regarding worker occupation, his/her weekly wage and the 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 are no unique correspondences between workers and task-types.[9] To avoid arbitrary assumptions, we thus perform the empirical analysis 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, and the tasks defined as *manual* belong in this group; (ii) FINGDEX, which is finger dexterity, and 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

[8] Author, Katz and Kearny (2006, 2008) and Lemieux (2006b) provide a full set of descriptive statistics on these data. According to Lemieux (2006b), the following are the drawbacks of the May/ORG CPS when used to analyze wage inequality: (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. Author, Katz and Kearny (2006, 2008) showed that the inclusion of data from the CPS March Issues help address some of these issues. In this paper, we follow the AKK strategy to build the database and we 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, such as answering phone calls, and cognitive tasks, such as 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.

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, and this group's tasks are typically defined as *cognitive* in nature; and (v) MATH, which refers to general education, analytical and mathematical skills and the ability to engage in problem solving, and it identifies the most typical *cognitive* tasks.[10] We adopt the five-group classification originally used in Author, Levy and Murnane (2003) instead of the three-group classification (manual, routine, cognitive) used in more recent literature, such as by Author and Dorn (2013), because our results and other side estimations show that the five-group classification is more effective than the three-group classification in identifying the dynamics of different types of tasks. In particular, we note that STS behaves differently from FINGDEX, and using a routine meta-group that consists of both would add noise to the data and obscure the results. The same occurs with DCP and MATH, whose remuneration levels appear to have both different dynamics 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. Therefore, growing wage inequality can be determined not only by higher skill premia but also by increases in the employment share of workers' groups with higher wage dispersion for reasons other than technology. For purposes of properly assessing variations in skill premia and their effects on wages, Lemieux suggested a reweighing kernel approach to simulate counterfactual changes in the wage distribution, while simultaneously maintaining the labor force composition constant and thereby isolating the so-called *price effect* from the *composition effect*.

Table 1 reports the results of the reweighing kernel analysis as applied to our data. In the first panel, we report overall, residual, and composition-adjusted wage variance for the initial and final sample periods, together with their changes in absolute terms. In the second panel, we consider the same statistics computed separately for above and below the median wage. Residual wage variance (RWI) 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, their cross-products

[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.

and the intensity of each group's occupational tasks (as defined in the previous section). The composition-adjusted residual wage variance is obtained by combining the actual price function with the composition function from the initial year.

Table 1: Wage variance for men 1986 – 2002

	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
Changes above and below the median wage			
below 50th	-0.0164	-0.0098	-0.0102
above 50th	0.0201	0.0228	0.0115

Wage variance is computed as the weighted variance among individual wages. Residual wage variance is computed as the weighted variance among individual residual wages, which are obtained regressing log hourly wages over worker's education, age and the intensity of each task-groups performed on duty. Composition-adjusted variance is then computed using the reweighing kernel approach described in Lemieux (2006) and using the initial period as base year.

The main conclusion that can be drawn from Table 1 is that adjusting for the composition of the labor force when computing RWI does not reduce the observed increase. As a matter of fact, the price effect appears even larger than the overall change in RWI, which suggests that changes in remunerations related to observable workers' characteristics played a non-negligible role in shaping the wage dynamics. In our intuition, the difference between our findings and those of Lemieux (2006) is due to the different definition of homogeneous cells used by Lemieux, which did not include the tasks performed by workers on duty (or workers' occupations). Given that the routinization hypothesis predicts that some jobs will become better-paying and that others will become lower-paying, it is reasonable to posit that the effects of remuneration average out if computed for cells in which all types of workers are included, which would bias the price effect downward. Note that this result does not hold when only wages above the median are considered in the analysis. As reported in the last line of Table 1, in this case the price effect accounts for no more than half of the overall change in RWI, thereby suggesting that the composition effect had a relevant role in shaping the wage dynamics in the right hand side of the distribution.

2.3 The influence of tasks on RWI

The following suggestive evidence is intended to understand whether occupational tasks are suitable candidates for explaining RWI. To this end, we regress the observed growth rate of RWI over the considered sample on task intensities. RWI is computed for narrowly defined groups of workers using the following characteristics: *education*, *experience*, *gender* and *occupations*. Because some of the 6,700 cells 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, RWI V_i is measured as the growth rate of the 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:

$$V_i = \alpha + \beta_1 ehf_i + \beta_2 fgx_i + \beta_3 sts_i + \beta_4 dcp_i + \beta_5 math_i + \varepsilon_i \tag{1}$$

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 (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 control variables that have been indicated in the literature as possible explanations for wage inequality, i.e., membership in unions, marital status and race [see Firpo, Fortin and Lemieux, 2011].[11]

[11] In detail, we first regress the overall growth rate of wages (first difference of log wages) on education, age dummies – used as a proxy for workers' experience – occupations, and their cross products. Next, we calculate the sum of the squared residuals obtained from the first-stage regressors for each cell V_i, which is eventually 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.

Table 2: Tasks effect on residual wage inequality growth by cells.

	Total		below 30th		30th-60th		above 60th	
	b/(se)	b/(se)	b/(se)	b/(se)	b/(se)	b/(se)	b/(se)	b/(se)
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

NOTES.

The results reported in Table 2 yield several insights. First, the estimations performed using the entire distribution reveal that (i) the non-routine analytic group has the only positive and significant coefficient among the tasks, (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 inequality among low-wage occupations, 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 appears to be significant – except for non-routine analytic tasks – which confirms that the bulk of the wage dynamics occurs at the periphery of the distribution. In general, previous results are broadly consistent with the predictions of the routinization hypothesis. Technology-complementary tasks (non-routine analytic) pushed wage inequality upward, whereas technology-substitute tasks (routine manual) pushed wage inequality downward. In the next section, we analyze in details this finding by means of a quantile regressions analysis.

2.4 Counterfactual Analysis

We now measure the change in wage inequality as the variation in wage gaps between selected percentiles 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 shown below, the following covariates are used: education, experience (proxied by age), task intensities by groups as defined in Section 2.1, union membership, marital-status and race. Following Firpo, Fortin and Lemieux (2011), we include the last three variables to control for factors that might affect wage inequality.[13] Tables 3 to 5 present the estimation outcome. For each percentile interval, the reported *price effect* (*composition effect*) represents the counterfactual 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

[12]This finding 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.

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 *price* of variable Z had changed.

From an aggregate perspective, the outcome of Table 3 mirrors 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, even when controlling for the composition effect, then the change in RWI is again positive in the upper echelon and negative in the bottom echelon (AKK). Note that the first line of Table 3 qualify AKK's finding by showing that the bulk of the wage dynamics is concentrated below the $30th$ and above the $60th$ percentiles, whereas wage inequality among the middle percentiles appears relatively stable. The last column of Table 4 also reveals that the largest contributor to the increase in overall wage inequality is the price effect between groups. Nevertheless, if the analysis were performed without distinguishing the between-group price effect from the within-group price effect, then the composition effect would appear as the largest contributor because the positive between-group price effect would be partially offset by the negative within-group price effect. This evidence reconciles our results with those found in Lemieux (2006). Table 4 clarifies several other features about the evolution of the price effect along the wage distribution. First, the reduction of inequality among lower percentiles is entirely due to a negative within-group price effect, which is partially compensated for by a positive composition effect. The between-group price effect is rather small and nonsignificant, suggesting that changes in wage differentials played a minor role in determining the overall reduction of wage dispersion among low-wage workers. Second, the analysis of percentiles around the median appears not to be insightful because the variation of inequality is nonsignificant. 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 the inequality variations in the middle percentiles are smaller than those at the periphery of the distribution. Finally, the price effect (strong) and the composition effect (mild) in the upper percentiles are both positive and jointly explain the large increase in the $60th - 95th$ wage gap.

Turning to the contributions of single variables, we find that the effect of *tasks* on within-group wage inequality outweighs the effects of all the other variables – including *education* and *experience*

– both when analyzing the distribution either as a whole or by echelons, and the only exception is the effect of education on top wages. Table 4 reveals that the largest contributors to the price effect within-group are the groups of *routine manual* and *non-routine analytic* tasks with symmetric and opposite effects on wage inequality. In other words, the first reduced inequality is located in the bottom percentiles, whereas the second increased inequality is located in the upper percentiles. In more detail, the variation in the remuneration of the group of routine manual tasks implied a reduction in the $10th-30th$ percentiles' distance 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 group of non-routine analytic tasks implied instead an increase of 4.2% in the $60th-95th$ percentiles' distance, which accounts for half of the variation in the upper echelon of the wage distribution. Table 5 demonstrates 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 ($+8\%$).

In general, the previous results support the job polarization theory's explanation of growing wage inequality. The effects of changes in the remuneration of occupational tasks is overwhelming in determining wage inequality and the signs of the effects are consistent with the predictions of RBTC. Moreover, wage differentials noticeably grew along the wage distribution and appeared to be the single most important source of changes in overall wage inequality. However, from the analysis performed separately on different echelons of the wage distribution, we learn that this result is fully driven by the effect of wage differentials in the upper echelon, whereas the main determinant of the evolution of wage inequality among lower wages is the contraction of wage dispersion *within* groups of workers. This piece of evidence appears to conflict with the notion of RBTC operating through a labor demand channel. For this reason, we next investigate an alternative formulation of RBTC based on the effects of technical change on the labor supply.

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

Table 3: Counterfactual Decomposition of changes in wage distances 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.

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 h_t, routine z_t, and manual l_t.

3.1 Labor demand and wages

The demand for labor is formed by a price-taker representative 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 – as opposed to the level – of technology, which thus implies 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

[14]Galor and Moav showed that a formulation in terms of growth rates helps disentangle the short- from the long-run effects of technological progress. In our partial equilibrium setup, such a formulation is mostly convenient

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 δs, which may direct the results. The demand of routine labor is thus 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] Thus, the occupation-specific wages are:

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

The wage rates (3)$-$(5) conform to the empirical evidence presented in the literature on the routinization hypothesis, in which abstract jobs represent the model counterparts of cognitive tasks.

for accommodating the different effects of technical change on labor demand and supply, as will shortly become clear.

[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.

Technology is assumed to push abstract wages upward, whereas its effect on routine and manual wages is undetermined although surely smaller than the effect on abstract wages.[16] As a consequence, technology 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 $[w_t^h/w_t^l = \beta(1-\delta g_t)^{-1}]$, whereas it 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 and isolates the labor supply effect, which is the central concern of our analysis.

3.2 Occupational Choices and the Labor Supply

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 the linear formulation suggested by Galor and Moav,

$$
\begin{aligned}
h_t^i &= a_i - (1-a_i)g_t && (6)\\
z_t^i &= 1 - (1-a_i)g_t && (7)\\
l_t^i &= 1 && (8)
\end{aligned}
$$

Equations (6) and (7) posit that the number of the efficiency units of abstract and routine labor increases with the ability level a_i – which replicates the assumption that ability reduces the cost of learning [Bartel and Sicherman, 1998 – and decreases with the pace of technical innovations g_t because we assume that, when there is technological progress ($g_t > 0$), workers must devote a fraction of their time to learning just 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 [Author,

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

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

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.

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 regarding wage inequality. In fact, any labor supply function entailing a learning advantage for abstract compared with routine workers would strengthen the migration of routine workers to abstract jobs when there are higher rates of technological progress, which would enhance the influence of technology on income inequality as a consequence. Second, only abstract occupations reward ability in a stationary technological environment, which 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 technology 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 a sorted mapping between occupations and income in equilibrium 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.

Figure 1: Income and Occupational Distributions

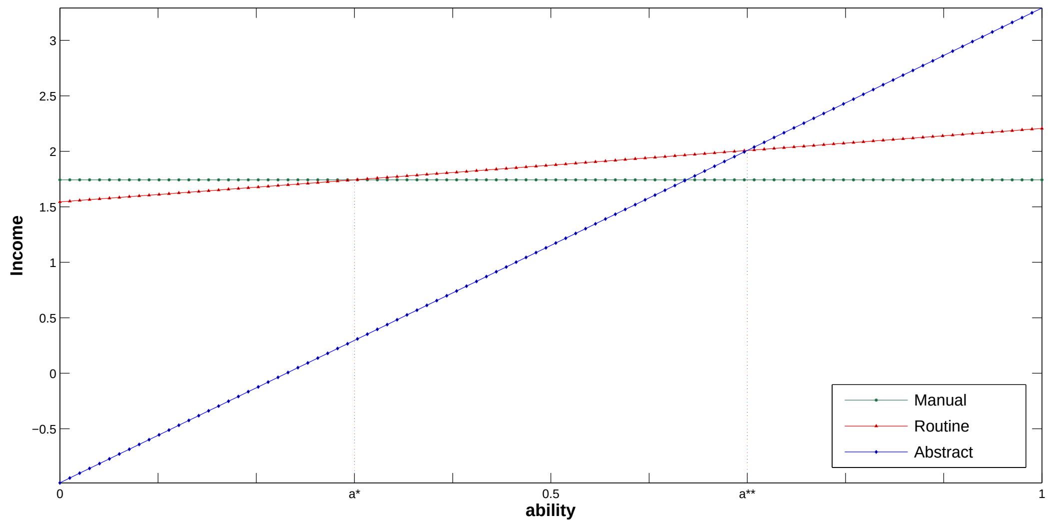

[17]Using equations (12) and (13), it can be shown that (i) there *always exists* a calibration of $\{\beta\,,\delta\}$ that guarantees a positive mass of individuals in each occupation, i.e., $0 < a^* \leq a^{**} < 1$, and (ii) it must fulfill two conditions, i.e., $\beta(1-g_t) > 1$ and $\delta < \frac{\beta-1}{g_t^2}$.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

3.3 Technical change and the income distribution

The equilibrium characterized in the previous section can be used to analyze the implications of technical change on the dynamics of the income distribution. Equations (9)−(11) imply that a variation of g_t shifts both the labor demand (wages) and supply (occupational choices), thus determining a new distribution of income in equilibrium and a new allocation of occupations along such new distribution. In particular, we show that workers' inflows generated by technical change imply variations in the ability intervals within occupations, 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 manual workers univocally grows.

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

Proposition 1 states that in presence of growing technical change wage differentials monotonically increase along the income distribution, whereas the effect on the composition of the labor force is non-monotonic, reducing employment in the middle echelon and increasing it at the extremes of the income distribution. This result replicates the job polarization dynamics suggested in the literature and is obtained in the model from 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 notable that downward migration only occurs in one direction because no manual worker finds the switching option desirable when the growth rate of technology increases. A higher level of g_t, in fact, has no effect on the relative wage between manual and routine jobs, although it reduces 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 would find the switching option desirable.

In the model, the highlighted composition effect is not the only channel through which technical change affects overall income inequality. In particular, it can be shown that workers' migration is not random but sorted across 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, thus affecting the within-group dispersion of wages. Moreover, because the downward migration is determined only by variations in the labor supply – the wage differential between manual and routine workers is fixed –, then the resulting within-group price effect operates entirely through a labor supply channel. In the following, we elaborate upon this point.

Proposition 2 (Skills distribution) *Consider the same economy defined in Proposition 1 and define the ability intervals of abstract, routine, and manual workers 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,*

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

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

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

It is clear that Proposition 2 has direct implications on within-group income inequality. When ability intervals increase (diminish), workers are less (more) homogeneous and income distribution within-group accordingly becomes less (more) equal. Proposition 2 therefore states that income inequality increases with the pace of technological progress for abstract and manual workers and decreases for routine workers. Such within-group price effects can be directly observed by defining the ratio of the highest to the lowest income in each occupational group, i.e.

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

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

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

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 instead null, but the result is due to the simplifying assumptions used in the liner model, which imply that manual income is constant without regard to the worker's ability. Note that the effect of g_t on within-group income inequality does not coincide with that on the ability intervals because the measure of wage inequality used is not dimensionless; therefore, variations in occupation-specific wages directly affect income inequality, although they have no direct effect on the ability intervals. In the cases of abstract and routine occupations, this mechanism operates on inequality in the same direction as that affecting the ability intervals, thereby univocally determining the overall effect. In the case of manual workers, the two mechanisms operate in opposite directions, thus making the overall effect of g_t *a priori* undetermined.[18] In general, from the results of Proposition 2 and the analysis of within-group income inequality, we can conclude that the effect of g_t operating though a labor supply channel closely mimics the one of RBTC observed in the data.

4 Conclusions

In this paper, we compare the effect of RBTC on the dynamics of inequality among low, medium and high wages. Depending on which echelon of the wage distribution is considered, two different patterns of wage inequality unfold in the analysis. In the upper echelon, technical change appears

[18] This consideration is always true, even though it is irrelevant in the linear model in which both the labor demand and the labor supply channels are shut down because within-group income inequality is always constant.

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

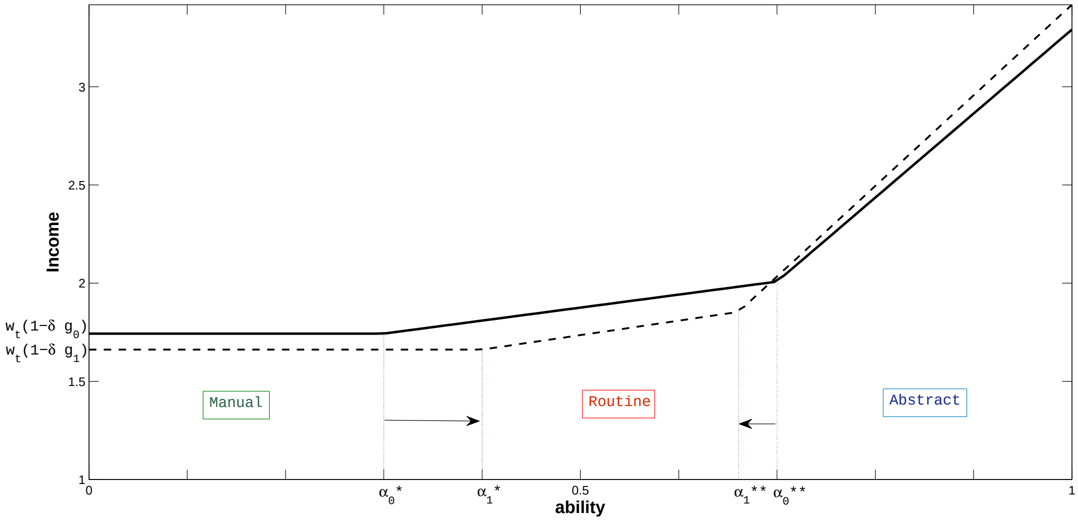

to operate on wage inequality through a labor demand channel. By raising the remunerations of technology-complementary tasks with respect to technology-substitute tasks, technical change pushes upward wage differentials among occupations comprising different occupational tasks, i.e., between-group wage inequality in our analysis, and the dispersion of wages within occupational groups, i.e., within-group wage inequality, where the latter effect is typically intended to be generated by increased remuneration to unobservable technology-complementary workers' skills. In the bottom echelon, we find instead that technology has no effect on wage differentials, even though it continues to have a strong effect on within-group wage inequality. We argue that this evidence points to a model in which RBTC operates through a labor supply channel. Specifically, 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 shows that there is no effect of technology on wage differentials in the lower echelon of the skill distribution and nonetheless argues that technical change has an effect on within-group inequality that mimics the effect observed in the actual data.

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, 2003, 118(4), pp. 1279-1333;

[5] 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;

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

[7] 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, 367-397;

[8] 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;

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

[10] Cortes, G. (2012): "Where have the routine workers gone? a study of polarization using panel data", mimeo University of Manchester;

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

[12] Firpo, S. and Fortin, N. and Lemieux, T. (2011): "Occupational tasks and changes in the wage structure", (No. 5542). Discussion paper series – Forschungsinstitut zur Zukunft der Arbeit;

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

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

[15] 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;

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

[17] Goos, M., Manning, A., and Salomons, A. (2011): "Explaining Job Polarization: The Roles of Technology, Offshoring and Institutions", Offshoring and Institutions (December 1, 2011);

[18] 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;

[19] Katz, L. F. and Murphy, K. M. (1992): "Changes in Relative Wages, 1963Ð1987: Sup- ply and Demand Factors", Quarterly Journal of Economics, 107(1), pp. 35Ð78;

[20] 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;

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

[22] 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), pp. 1333Ð81;

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

[24] 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 (November);

[25] Piketty, T., and Saez, E. (2006): "The Evolution of Top Incomes: A Historical and International Perspective", The American Economic Review, 200-205;

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

A The Counterfactual Quantile Regressions Approach

To perform the wage inequality decomposition presented in Section 2.4, we employ the extension proposed by AKK of Machado and Mata (2005) CQR. Hereafter we summarize the steps followed. 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. 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, and for each draw repeat the procedure in the initial and final period. Collect the estimated coefficients $\hat{\beta}_{\theta_i,t}$ into a 10.000 rows vector $\hat{B}_t$. Then, the marginal density of w_t is obtained by generating a random sampling $x^*_{i,t}$ from the rows of X_t and obtaining the unconditional distribution of wages as:

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

Finally, compute the simulated unconditional quantile $\hat{\theta}$ as $\hat{Q}_{\theta,t}(w^*_{i,t})$.

In general, wage inequality growth 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 $0 \leq a^* < a^{**} \leq 1$, item (i) directly follows from the fact that $\frac{\partial a^{**}}{\partial g_t} < 0$ and item (ii) from the fact that $\frac{\partial a^{*}}{\partial g_t} > 0$. Then, item (iii) 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, i.e.

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

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

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

Item (i) and (ii) then follows respectively from $\frac{\partial \overline{a}^l}{\partial g_t} > 0$ and $\frac{\partial \overline{a}^z}{\partial g_t} < 0$. Regarding item (iii), it can be shown that

$$\frac{\partial \overline{a}^h}{\partial g_t} = \frac{\delta - \delta^2 g_t^2}{(1 + \delta^2 g_t^2)}$$

and therefore condition $g_t < \sqrt{\delta^{-1}}$ guarantees that $\frac{\partial \overline{a}^h}{\partial g_t} > 0$.

Table 4: Counterfactual Decomposition of changes in wage distances by quantiles

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.

Table 5: Counterfactual Decomposition of changes in wage distances by quantiles

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.