Occupational Tasks and Changes in the Wage Structure [*]

By Sergio Firpo, Nicole M. Fortin, and Thomas Lemieux

*This paper argues that changes in the returns to occupational tasks linked to offshorability and the technological content of work have contributed to changes in the wage distribution over the last three decades. Occupational tasks are measured using the O*NET data set, and wage data are from the Current Population Survey (CPS). Using a decomposition procedure based on RIF-Regressions, we find that technological change and de-unionization played a central role in changes in the wage distribution during the 1980s and 1990s, while offshorability became an important factor from the 1990s onwards.*

JEL: J3, J5

Keywords: wage inequality, polarization, occupational tasks, offshoring, RIF-regressions.

[*] Firpo: São Paulo School of Economics, Fundação Getulio Vargas, sergio.firpo@fgv.br; Fortin and Lemieux: Vancouver School of Economics, University of British Columbia, nicole.fortin@ubc.ca and thomas.lemieux@ubc.ca. We would like to thank Matias Cortes for providing supporting PSID analysis. We are indebted to two anonymous referees, David Autor, David Card, David Dorn, Martin Goos, Caroline Hoxby, Luigi Pistaferri, Anna Salomon, Alexandra Spitz-Oener, Frank Vella, and seminar participants at UC Berkeley, Calgary, CEMFI, Cornell, Dalhousie, Georgetown, UCL, LSE, HEC-Montreal, McMaster, Montreal, Ottawa, Toronto, Calgary, Hawaii, Paris II, PSE, Columbia, Stanford, Victoria, Collegio Carlo Alberto, UC Davis, the 2010 NBER Summer Institute, the T.A.S.K.S. International Workshop in Nuremberg, the 2010 EALE/SOLE Conference, and the ACAM Conference in Tokyo for useful comments on this and earlier versions of the manuscript. Financial support from CNPq-Brazil for Firpo, from CIFAR and SSHRC for Fortin, and from SSHRC for Lemieux was greatly appreciated. The usual disclaimer applies.

Traditionally, most studies of changes in inequality and the wage structure have focused on explanations based on changes in the returns to skills, like education and experience (e.g. Katz and Murphy, 1992) or institutions (e.g. DiNardo, Fortin, and Lemieux, 1996).[1] Recently, more attention has been paid to the potential role of occupations in changes in wage inequality. This shift happened for several reasons.

First, Autor, Levy, and Murnane (2003), Goos and Manning (2007), and Autor, Katz and Kearney (2006) have proposed a new explanation for changes in wage inequality based on a more "nuanced" view of skill-biased technological change. The idea is that the introduction of computer and information technologies has not simply depressed the relative demand for less skilled workers, as was assumed in early studies such as Berman, Bound, and Griliches (1994). Rather, computer and information technologies have depressed the return to "routine" tasks that can now be executed by these technologies. Autor, Katz and Kearney (2006), Goos and Manning (2007), Autor and Dorn (2013), and Michaels, Natraj, and Van Reenen (2010) argue that this nuanced view of technological change can help account for the polarization of wages that has been observed since the late 1980s. Under this type of technological change, it is plausible that moderately skilled workers who used to perform routine tasks experienced a decline in relative wages during this period. Technological change could thus explain why wages in the middle of the distribution fell more than wages at the bottom and top end of the distribution.[2]

This more nuanced view of technological change puts occupations at the forefront of the inequality debate since the task content of work (routine nature of the job, cognitive skills required, etc.) is typically determined at the occupational level.[3] Occupations are, therefore, a key empirical channel through which we can assess how technological change affects the wage structure. An important empirical implication of this more nuanced view of technological change, that we discuss below, is that changes in the wage structure within and between occupations should be systematically related to the type of tasks performed in these occupations.

A second reason for looking at the contribution of occupations to changes in the wage structure is offshoring. Early explorations of the role of international trade in changes in inequality have focused on the role of trade in final products, defined at the industry level. It was later argued (e.g.

[1]The role of industrial change due to de-industrialisation and foreign competition was also explored in some of the early studies such as Murphy and Welch (1991), Bound and Johnson (1992), and Freeman (1995).

[2]Acemoglu and Autor (2011) develop a formal model to show how this could happen in a model with three skill levels (high, middle, and low).

[3]Most studies have either used data from the Dictionary of Occupation Titles (DOT) or the more recent Occupational Information Network (O*NET) to get information about the task content of jobs. Since jobs are defined on the basis of a detailed occupational classification, this naturally lead to an analysis at the occupational level.

Feenstra and Hanson, 2003) that trade in intermediate inputs was a more promising explanation than trade in final goods and services. More recently, Grossman and Rossi-Hansberg (2008) have proposed a model of the global production process where tasks are tradeable. In their model, reductions in the cost of offshoring tasks have effects much like factor-augmenting technological progress, boosting the productivity of workers whose tasks become easier to move offshore and thereby rising their wages. As in the case of technological change, occupations are the key channel through which offshoring can contribute to changes in wage inequality.[4]

Although occupations now feature prominently as a possible channel for recent changes in wage inequality, the role of occupations in these changes has not been systematically investigated yet. Some studies do suggest an important role for occupation-based explanations. Goos and Manning (2007) show that the composition effect linked to changes in the distribution of occupations accounts for a substantial part of the increase in inequality in the United Kingdom. Autor, Katz and Kearney (2008) provide evidence that, consistent with a nuanced view of technological change, the share of employment in occupations in the middle of the wage distribution has declined over time. Using a spatial equilibrium approach, Autor and Dorn (2013) show that local labor markets that are more specialized in routine jobs experienced more polarization. While these findings suggest a potentially important role for occupations, it remains to be seen how much of the total change in the distribution of wages can precisely be accounted for by occupation-based explanations by comparison with other changes, such as de-unionization or increasing returns to education.

The goal of this paper is to fill this gap by systematically investigating the contribution of occupations to changes in the distribution of wages in the United States. We do so by first introducing a wage setting model to clarify the connection between skills, tasks, and wages. The model closely follows the approach of Acemoglu and Autor (2011) where skills are used to produce tasks in an occupation, and wages depend on both the market price of tasks, and the amount of tasks produced by the worker. In this setting, conditional on skills, changes in occupational wages depend on changes in task prices. Although these tasks prices are not directly observed, we use detailed data from the Occupational Information Network (O*NET) to construct a set of task content measures that are then used to predict changes in task prices linked to offshoring and technological change. For example, as in Autor, Levy, and Murnane (2003), jobs where work is highly repetitive get a high score for the "automation/routine" work content variable. Since

[4]Ebenstein, Harrison, McMillan and Phillips (2013) look at the impact of offshoring on wages at the occupation level. They find mixed results depending on whether jobs are offshored to low or high wage countries.

these jobs can be more easily replaced by computer-operated machinery, we expect task prices, and thus wages in these jobs tend to decline over time relative to jobs with a low score for the "automation/routine" work content variable. Similarly, we use several task content measures to capture the potential offshorability of jobs, as in Blinder (2009) and Jensen and Kletzer (2010)).

Using Current Population Survey (CPS) data for 1976 to 2012, we then quantify the contribution of changes in task prices and other factors to changes in the distribution of wages over that period. We do so using a decomposition method based on the recentered influence function regression approach of Firpo, Fortin, and Lemieux (2009). This approach enables us to evaluate the contribution of changes in task prices compared to other explanations such as de-unionization and changes in the labor market-wide returns to general skills (labor market experience and education). We find that technological change and de-unionization played a relatively central role in the 1980s and 1990s, but had little effect in the 2000s. Market-wide increasing returns to education played an important role in all three decades, while offshorability became an important factor in the 1990s and 2000s.

The paper is organized as follows. In Section I, we present a wage setting model that helps frame the empirical analysis by clarifying the connection between skills, tasks, and wages. We present some descriptive evidence supporting the model and discuss its connection with the decomposition procedure used for the main empirical analysis. Section II describes the wage data used, introduces the measures of task content computed from the O*NET data, and explains how they are linked to the concepts of technological change and offshorability. Section III provides a summary of the decomposition methodology based on recentered influence function regressions. The main decomposition results are presented in Section IV. We conclude in Section V.

I. Wage Setting in Occupations

A. Skills, Tasks, and Wages

Most of the wage inequality literature follows a traditional Mincerian approach where wages are solely determined on the basis of (observed and unobserved) skills. Equilibrium skill prices depend on supply and demand factors that shape the evolution of the wage structure over time. Underlying changes in demand linked to factors like technological change and offshoring can certainly have an impact on the allocation of labor across industry and occupations, but ultimately wage changes are only linked to changes in the pricing of skills. Acemoglu and Autor (2011) refer to this approach as the "canonical model" that has been used in many influential studies, such as Katz and Murphy (1992).

There is increasing evidence that the canonical model does not provide a satisfactory explanation for several important features of the evolution of the wage structure observed over the last few decades. This is discussed in detail in Acemoglu and Autor (2011) who mention, among other things, two important shortcomings of the canonical model. First, it cannot account for differential changes in inequality in different parts of the distribution, such as the "polarization" of the wage distribution of the 1990s. Second, the model does not provide insight on the contribution of occupations to changes in the wage structure because it does not draw any distinction between "skills" and "tasks". Acemoglu and Autor (2011) address these shortcomings by proposing a Ricardian model of the labor market where workers use their skills to produce tasks, and get systematically allocated to occupations (i.e. tasks) on the basis of comparative advantage.[5]

We closely follow Acemoglu and Autor (2011) in the way we introduce the distinction between skills and tasks in our wage setting model. Unlike Acemoglu and Autor (2011), however, we do not attempt to solve the full model of skills, tasks, and wages by modelling how workers choose occupations, and how supply and demand shocks affect wages in general equilibrium. One advantage of our partial equilibrium approach is that we don't have to impose restrictive assumptions to help solve the model. For instance, Acemoglu and Autor (2011) have to work with only three skill groups (but many occupations/tasks) to get interesting predictions out of their model. As a result, the law of one price holds within each skill group in the sense that wages are equalized across occupations, conditional on skill. This is a strong prediction that is not supported in the data, and that we can relax by allowing for a large number of skill categories.[6] This limits our ability to solve the model in general equilibrium, which is beyond the scope of this paper. Yet, the fact that workers systematically sort into different occupations/tasks on the basis of their skills has potentially important implications for the interpretation of our results. We discuss these issues in more detail at the end of this section.

Like Acemoglu and Autor (2011), we assume that an occupation j involves producing a task or occupation-specific output Y_j which is one input in the firm's production function. But instead of just working with three skill types, we assume that workers are characterized by a k-dimension set of skills $S_i = [S_{i1}, S_{i2}, ..., S_{iK}]$. Some of these skills (like education and experience) are

[5]Note that since different tasks are being performed in different occupations, we can think of these two concepts interchangeably.

[6]See, for instance, Heckman and Scheinkman (1987) and Gibbons et al. (2005) for evidence of occupational wage differences among workers with similar observed and unobserved productive characteristics.

observed by the econometrician, others (like ability and motivation) are not. The amount of occupation-specific task Y_{ij} produced by worker i in occupation j is assumed to linearly depend on skill:

$$Y_{ij} = \sum_{k=1}^{K} \alpha_{jk} S_{ik}, \tag{1}$$

where the productivity of skills α_{jk} are specific to occupation j. Firms then combine tasks to produce final goods and services according to the production function $Q = F(Y_1, ..., Y_J)$ where Y_j (for $j = 1, .., J$) is the total amount of (occupation-specific) tasks produced by all workers i allocated to occupation j.[7]

Under the assumption that wages are set competitively, workers are paid for the value of tasks they produce. Worker i who produces Y_{ij} units of occupation-specific task j is thus paid a wage of $p_{jt}Y_{ij}$, where p_{jt} is the market price of each unit of task Y_{ij} produced at time t. We also allow wages to depend on year and occupation specific factors δ_t and c_j, where δ_t could capture, for instance, general productivity shocks, while c_j could be thought as reflecting compensating wage differentials. In the empirical analysis, we also consider other factors Z_{it} such as institutions (e.g. union status) and discrimination (e.g. race and gender) that affect wages in a way that is unrelated to task output. This yields the wage equation:

$$w_{ijt} = \delta_t + Z_{it}\psi_t + c_j + p_{jt}Y_{ij} \equiv \delta_t + c_j + Z_{it}\psi_t + p_{jt} \sum_{k=1}^{K} \alpha_{jk} S_{ik}. \tag{2}$$

As in Acemoglu and Autor (2011), a critical assumption embedded into equation (2) is that the mapping of skills into tasks (the parameters α_{jk} in the wage equation) does not change over time, while task prices p_{jt} are allowed to change over time. This means that, in this model, the effect of demand factors such as offshoring and technological change solely goes through changes in task prices. In this setting, technological change and offshoring provide a way for firms of producing the same tasks at a lower price. Take, for instance, the case of call center operators who use their skills to produce consumer service tasks (check customer accounts, provide information about products, etc.). When these tasks are simple, like providing one's balance on a credit card, the call center operators can be replaced by computers now that voice recognition technology is advanced enough. In the case of more complex tasks such as IT support, computers are not

[7]This specification is also closely related to the "skill-weights" approach of Lazear (2009) where different jobs require the use of different linear combinations of skills.

sophisticated enough to deal with customers but these tasks can now be offshored to lower paid workers in India. In these examples, the quantity of task produced by call center operators of a given skill level does not change, but the wage associated with these tasks changes in response to technological change and offshoring. At the limit, if the task price in an occupation becomes low enough the occupation will simply disappear, which is the way Acemoglu and Autor (2011) model the impact of "routine-biased" technological change.

In other cases the assumption that the mapping between skills and tasks is constant over time may be unrealistic. For instance, in highly technical or professional occupations where cognitive skills are important for producing tasks, advances in computing likely enable workers with a given set of skills to produce more tasks than they used to. In that setting, when wages increase for these workers, equation (2) would suggest that task prices have increased, while the underlying explanation may instead be productivity changes linked to changes in the α_{jk}'s. Since p_{jt} and α_{jk} enter multiplicatively in equation (2), it is not possible to empirically distinguish the impact of changes in these two factors. Ultimately, the product of p_{jt} and α_{jk} is an occupation-specific return to skill at time t, and the main goal of the paper is to quantify the contribution of changes in the these occupation-specific returns to skill on changes in the wage distribution, controlling for other factors usually considered in the inequality literature. For the sake of simplicity we interpret these changes in returns as changes in task prices, but acknowledge that they could also reflect occupation-specific productivity effects.

When task prices are allowed to vary across occupations in a completely unrestricted way, it is difficult to interpret the contribution of changes in task prices to changes in inequality in an economically meaningful way. Following Yamaguchi (2012), we assume that task prices are systematically linked to a limited number of task content measures available in data sets like the Dictionary of Occupational Titles or the O*NET. The idea is that two different occupations where the task content measure for, say, "routine work" is the same will be equally affected by "routine-biased" technological change. In the empirical part of the paper we use a set of five task content measures from the O*NET that are described in detail in the next section. We use the following linear specification for task prices:

$$p_{jt} = \pi_{0t} + \sum_{h=1}^{5} \pi_{ht} T_{jh} + \mu_{jt}, \tag{3}$$

where T_{jh} are the task content measures. These task content measures are assumed to be time invariant for two reasons. First, it has proven difficult to construct consistent measures of the task content of occupations over time because of data limitations (see, e.g., Autor, 2013). More importantly, we use the task content measures as an economically interpretable way of reducing the dimension of the occupational space. Results would be hard to interpret if the way in which task content characterized occupations was also changing over time.[8]

Since the T_{jh}'s do not change over time, changes in task prices p_{jt} are solely due to change in the parameters π in equation (3). These parameters can be interpreted as the returns to task content measures T_{jh} in the task pricing equations. Relative to the existing inequality literature, the main contribution of our decomposition exercise is to look at the contribution of changes in these returns to task content measures (changes in the π_{ht}'s) on changes in the wage distribution, in addition to more standard explanations explored in the literature.

The effect of changes in π_{ht} on changes in the wage distribution are complex. To see this, consider the wage equation obtained by substituting equation (3) into (2):

$$w_{ijt} = \delta_t + c_j + Z_{it}\psi_t + \left[\pi_{0t} + \sum_{h=1}^{5} \pi_{ht} T_{jh} + \mu_{jt}\right] \sum_{k=1}^{K} \alpha_{jk} S_{ik}. \tag{4}$$

Since task prices and skills enter multiplicatively into the wage equation, a change in task prices linked to changes in the π_{ht} parameters has an impact on both the between- and within-group dimensions of inequality. For instance, even if the α_{jk} parameters were the same in all tasks/occupations, changes in π_{ht} would increase wage dispersion between occupations as long as average skills (e.g. education, one of the elements of the skill vector S_i) varied across occupations. Furthermore, since some dimensions of skills are unobserved, changes in π_{ht} also affect within-occupation inequality even after controlling for observable skills like education and experience. More generally, the impact of changes in π_{ht} (or other factors) may be quite different at different points of the distribution, depending on the distribution of T_{jh} and S_{ik}.

Note also that the intercept term π_{0t} captures changes in skill prices that are common to all occupations. From that point of view, equation (4) allows for more standard increases in returns to skills that are not linked to occupational tasks. This provide a rationale for looking at the contribution of market-wide changes in the returns to education or experience in our

[8]Note that Yamaguchi assumes that the parameters α_{jk} are also functions of the task content variables T_{jh}, something we do not do since we would then need to be more specific about the way we introduce the K observed and unobserved skill components (corresponding of each parameter α_{jk}). More importantly, the question of whether or not the T_{jh}'s should be allowed to change over time in this setting is just a more structured way of thinking about the implications of possible changes in α_{jk}, an issue that we have already discussed.

decomposition exercise.

In principle, one could treat the wage setting equation (4) as a structural model, make a number of assumptions about the distribution of the unobserved components in the skill vector S_i, and estimate the parameters of the model. We use a simpler and less parametric approach by carrying out a decomposition where wage changes at each point of the distribution are decomposed in a number of components linked to changes in the distribution of observed covariates (the observed skill components and the task content measures), and in the returns to these covariates.

Our decomposition approach is explained in detail in Section III. It can be viewed as a generalization of the familiar Oaxaca-Blinder decomposition as it identifies the contribution of "price" (e.g. π_{ht}) and "quantity" (e.g. T_{jh} and S_{ik}) effects to changes in the wage distribution. The close connection between changes in the π_{ht} parameters and the "price", or wage structure, component of the decomposition is shown explicitly in Appendix B. There, we also discuss a simplified example to help establish the connection.

In addition to carrying the full decomposition with detailed occupations (3-digits) later in the paper, we also look explicitly at the connection between the task content measures and changes in the between- and within-occupation wage dispersion at a coarser occupation level (2-digits) in Appendix C. There are large differences in the changes in the level and dispersion of wages across occupations. This is illustrated in Figure 2 in the case of men over the 1990s. The figure shows the change in wages by decile (as a function of base period wages) in three broad occupation groups: food workers, skilled production workers, and engineers. In some "middle-end" occupations like production workers, all wage deciles decline in real terms, while they tend to increase in other occupations at the top-end (e.g. engineers) or low-end (e.g. food workers) of the distribution. Furthermore, wage dispersion increases for engineers (top wage deciles increase more than lower wage deciles) while the opposite happens for food workers (production workers are more neutral in this regard).

Appendix C explores these issues in more details using changes in average occupational wages as a summary measure for the level of wages, and the standard deviation as a summary measure for (within-occupation) wage dispersion. We reach two important conclusions in that analysis. First, we find that changes in these two summary measures are closely connected as occupations experiencing relative gains in average wages also tend to experience growth in wage dispersion. The correlation between changes in the mean and standard deviation across occupations is large and positive (0.44), and increases to 0.57 when only looking at non-agricultural occupations. This is consistent with both of these dimensions of wage dispersion depending on the same underlying factor (task prices p_{jt}) in the wage setting model.

Second, we find that most of the changes in both the level and dispersion of occupational wages can be explained by the task content measures constructed using the O*NET data. The adjusted R-square of a regression of changes in average wages (or changes in standard deviations) on the five tasks content measures is large and positive (0.52 in both cases). These results suggest that changes in the mean and standard deviation are in a large part driven by the same underlying changes in task prices (p_{jt}), and that these changes in task prices are well predicted by our task content measures (the T_{jh}'s). Taken together, these results suggest that the simple wage setting model in equation (4) is rich enough to capture the main features of the changes in the distribution of wages across occupations.

B. Selection and general equilibrium issues

One important feature of our wage setting model is that workers with a given skill vector S_i earn different wages in different occupations j. By contrast, in Acemoglu and Autor (2011) the law of one price holds (same wage in different occupations) within each of the three skill groups they consider. The source of this difference is that Acemoglu and Autor's model corresponds to the case considered by Rosen (1978) where the number of tasks largely exceeds the number of worker types. Rosen also considers the opposite case where the number of worker types is much larger than the number of tasks. He shows that wages (and return to skill more generally) depend on tasks/occupations in that alternative setting. This provides the theoretical rationale for our wage setting model. As we discuss above, one consequence of using more skill types than tasks is that it is difficult to derive a closed form for the model.[9]

Another consequence is that, as in a Roy model, workers will systematically sort themselves into different occupations on the basis of comparative advantage. For example, occupations where cognitive skills are important for producing tasks will attract workers with higher cognitive skills. More importantly, when task prices p_{jt} change in response to demand shocks (technology, offshoring, etc.), workers will tend to leave occupations where p_{jt} has declined to get a higher wage in another occupation. If relatively more (or less) skilled workers are more likely to change occupation in response to a decline in p_{jt}, this compositional change will tend to confound the effect of changing p_{jt}.

[9]Jung and Mercenier (2010) use a similar strategy to look at the effect of offshoring and technological change on the wage distribution, but impose fairly restrictive assumptions on the technology. See also Cortes (2012).

Since we only use cross-sectional data from the CPS in our empirical analysis, we can only control for selection based on observables, i.e. adjust for the observable components of the skill vector S_i. But unless the distribution of unobservable skills within occupations remains constant conditional on observable skills (ignorability assumption), the effect of changes in p_{jt} will still be confounded by changes in the distribution of unobservable skills within occupations.

The results of a recent paper by Cortes (2012) can be used to assess the importance of potential biases linked to sorting on the basis of unobservables. Cortes controls for selection (on unobservables) using longitudinal data from the Panel Study of Income Dynamics (PSID) where he allows workers to have an occupation-specific match term (or fixed effect) for three groups of occupations: routine, non-routine manual, and non-routine cognitive occupations. He concludes that the polarization phenomenon where wages in routine occupations decline relative to the two other job types remains after controlling for unobservables.

In Appendix Figure A1, we reproduce the selection-adjusted changes in occupational wages estimated by Cortes (2012), and compare them to those that would be obtained controlling only for observable covariates available in the CPS.[10] Relative to the base occupation (non-routine manual), the figure shows that wages in routine occupations plummeted starting in the mid-1980s, while wages on non-routine cognitive jobs (those at the top end) increased rapidly. Interestingly, the results indicate that controlling for selection based on unobservables makes the wage changes even more dramatic. This suggests that, if anything, using the CPS leads to an understatement of the changes in wages, and thus in task prices, across occupations. We conclude from this exercise that selection based on unobservables is unlikely to overstate the contribution of changes in task prices in our decomposition exercise.

General equilibrium considerations would also likely decrease the contribution of changes in task prices to changes in inequality. When an occupation is hit negatively by a demand shock, task prices decline and workers move to other occupations, thereby reducing wages in these other occupations. In the model of Acemoglu and Autor, these supply adjustments would eventually lead to a situation where wages across occupations would be equalized again within a given skill group. Wages in the skill group initially allocated to the adversely affected occupations (the "middle" skill group in their analysis of the impact of routine-biased technological change) would decline, but this would be captured in a decomposition by conventional changes in returns to skill.

[10]We are very grateful to Matias Cortes for providing the estimates shown in the figure.

In summary, overlooking both general equilibrium effects (in theory) and selection effects (in the PSID data) likely understate the contribution of occupation-specific demand shocks (captured by our task content variables) to changes in the wage distribution. Thus, if we find that the task content of occupations helps account for some of the change in the wage distribution in our decomposition exercise, we can be confident that this truly reflects the contribution of occupation-specific demand shocks.

II. Data

A. Wage Data

The empirical analysis is based on data for men and women, studied separately, from the Outgoing Rotation Group (ORG) Supplements of the CPS.[11] The wage measure used is an hourly wage measure deflated to 1979 real dollars using monthly CPI. For workers paid by the hour, we use a direct measure of the hourly wage rate. For workers not paid by the hour, the hourly wage rate is computed by dividing earnings by hours of work. CPS weights are used throughout the empirical analysis. We pool several years of data together to improve the precision of the estimates. For the first period (1976-78 to 1988-90), we start with data from the May CPS for the years 1976 to 1978.[12] For the second, and main period of analysis, we use 1988-90 as the base year and 2000-02 as the end year to make sure we fully capture all the changes that occurred during the 1990s. Our final period goes from 2000-02 to 2011-12.[13]

We consider changes in men and women's wage distributions separately given the substantial amount of occupational segregation that persists to this day. The substantial overrepresentation of women in pink collar jobs and of men in blue collar jobs leads to different task content by gender.[14] Table 1 shows that by 2000-02 women have overtaken men in their relative representation among professionals and technicians, but women are still overrepresented by 20 percentage points

[11]The data files were processed as in Lemieux (2006b) who provides detailed information on the relevant data issues. Sample means are provided in Appendix Table A1.

[12]The reason we use the May CPS instead of the MORG CPS for 1979 or 1980 is that union status was not asked in the MORG CPS until 1983. Since inequality was relatively stable during the 1970s (see DiNardo, Fortin, and Lemieux, 1996), the precise choice of base year for studying changes in inequality during the 1980s should not have much impact on the results.

[13]We note that there was a dramatic change in the coding of occupations when the 2000 census classification was introduced in 2003, followed by smaller changes in 2010. We do not attempt to crosswalk the pre-2000 occupation codes to the 2010 codes. Rather we crosswalk the more numerous O*NET occupations codes with both the 1980s and 1990s codes, on the one hand, and with new post-2000 codes, on the other hand. There was also a substantial change in the coding of occupations when the 1980 census classification was introduced in 1983. For this change, we use a crosswalk to keep a reasonably consistent definition of occupations between 1976-78 and 1988-90.

[14]In addition, the 1980s was a decade, unlike the two subsequent ones, that continued to see large increases in female labor force participation, likely involving changes in selection into the labor market (see Mulligan and Rubinstein, 2008). Therefore, we de-emphasize the results for women prior to the 1990s.

among clerical and sales workers, and underrepresented among primary sector, construction, and transportation workers.

Consistent with Autor, Katz and Kearney (2006), Figures 1a and 1b show that changes in real wages at each percentile of the wage distribution follow a U-shaped curve, among both men and women, from 1988-90 to 2000-02. In the figure, we also contrast these wage changes with those that occurred before (1976-78 to 1988-90) which were largely monotonic, and those that followed (2000-02 to 2011-12) which were more J-shaped. The figure illustrates that wages at the very top have generally increased much more than wages in the middle of the distribution, resulting in increased top-end inequality. By contrast, inequality in the lower half of the distribution increased rapidly during the 1980s, but decreased sharply after 1988-90 as wages at the bottom grew substantially more than those in the middle of the distribution.

The very bottom part of the wage distributions (below the 10th centile) has seen some gains since 2000-02. This is somewhat surprising since recessions are typically believed to have a particularly negative impact at the bottom end of the distribution. For women, the large swings at the bottom of the distribution, especially in the 1980s, have been attributed to minimum wage effects (DiNardo, Fortin, and Lemieux, 1996; Lee, 1999; Autor, Manning and Smith, 2010). More generally, wage changes between 2000-02 and 2011-2012 may be partly driven by differences in macroeconomic conditions and composition effects since the unemployment rate in the aftermath of the Great Recession was still unusually high.[15] As it turns out, however, inequality in hourly wages does not exhibit much of a cyclical pattern. This can be seen in the case of the 50-10 and 90-50 log wage differentials for both men and women in Appendix Figure A2. This suggests that macroeconomic conditions likely play little role in the key inequality changes documented in this paper.

B. Occupational Measures of Technological Change and Offshoring Potential

Like many recent papers (Goos and Manning (2007), Goos, Manning and Salomons (2010), Crinó (2010)) that study the task content of jobs, and in particular their offshorability, we use the O*NET data to compute our measures of technological change and offshoring potential.[16] The construction of task content indexes has generally followed two alternative paths, either a top-down approach where economic reasoning guides the choice of job characteristics of interest

[15] By contrast, the overall state of the labor market was more or less comparable in the other years considered in the analysis. The average unemployment rate for the 1976-78, 1988-90, and 2000-02 period is 6.2, 5.9, and 4.8 percent, respectively, compared to 8.1 percent for 2011-12.

[16] We use the O*NET 13.0 available from National Center for O*NET Development.

(as in Autor, Levy, and Murname (2003)), or a bottom-up approach, where a principal components analysis is used to construct less interpretable, but orthogonal, principal components (as in Poletaev and Robinson (2008)).[17] We follow the first approach because we are trying to investigate the relatively recent phenomenon of offshoring on wages, rather than using a take-no-stand statistical description of the task content of jobs.

The construction of our index of potential offshorability follows the pioneering work of Jensen and Kletzer (2010) [JK hereinafter], while incorporating some of the criticisms of Blinder (2009). Our aim is to produce indexes for all 3-digit occupations available in the CPS. In the spirit of Autor, Levy, and Murnane (2003), who used the Dictionary of Occupational Titles (DOT) to measure the routine vs. non-routine, and cognitive vs. non-cognitive aspects of occupations, JK use the information available in the O*NET, the successor of the DOT, to construct their measures. The O*NET content model organizes the job information into a structured system of six major categories: worker characteristics, worker requirements, experience requirements, occupational requirements, labor market characteristics, and occupation-specific information.

Like JK, we focus on the "occupational requirements" of occupations and also add some "work context" measures to enrich the "generalized work activities" measures. JK consider eleven measures of "generalized work activities", subdivided into five categories: information content, internet-enabled, face-to-face contact, routine or creative nature of work, "on-site" nature of work. Blinder (2009) argues that a difficulty with objective indexes of non-offshorability is incorporating two important criteria: a) that a job needs to be performed at a specific U.S. location, and b) that the job requires face-to-face personal interactions with consumers. We thus pay particular attention to the "face-to-face" and "on-site" categories in the construction of our indexes. Indeed, Autor and Dorn (2013) argue that only these two components should be included as job characteristics of non-offshorability. As shown below, here no-decision making helps differentiate service occupations from professional occupations. We also note that while JK see information content and decision-making as characteristics of non-offshorability, Autor et al. (2003) associate decision, control and planning (DCP) with technological change. Caroli and Van Reenen (2001) further argue that change in decision-making is associated with skill-biased organizational change. These differing views reflect the fact that technological change, organizational change, and offshoring are changes in work production that most often do not take place in isolation. Thus, it may be an heroic attempt to parse them out completely.[18]

[17] Autor (2013) discusses the trade-off involved in the alternative approaches.

[18] The correlations between our task content measures are quite low, generally below 0.15 using the data from Table 1.

We consider five task measures similar to JK, but think of our first two measures, *i)* the information content of jobs, and *ii)* the degree of automation of the job and whether it represents routine tasks, as being more closely linked to technological change, including advances in computerization, information and communication technologies (ICT). Thus, we refer to this group of task measures as "Technological Change". These two measures aggregate the following task subcomponents *i)* information content: getting information, processing information, analyzing data or information, interacting with computers, and documenting/recording information; and *ii)* automation/routine: degree of automation importance of repeating same tasks, structured versus unstructured work, pace determined by speed of equipment, spend time making repetitive motions.[19] Our first measure, "information content", regroups JK's first two categories; it identifies occupations with high information content that are likely to be affected by ICT technologies. Our second measure, "automation/routinization", is constructed using work context measures to reflect the degree of potential automation/robotization of jobs and is an update on the manual routine index of Autor et al. (2003). Like Blinder (2009) we acknowledge that there is some degree of overlap with offshorability, but more so in the 2000s where internet-enabled jobs with high information content could be offshored if there are no mitigating factor, such as involving strategic decision-making.

Our three remaining task measures, *iii)* the importance of face-to-face contact, *iv)* the need for on-site work, and *v)* the importance of decision making on the job, are meant to capture features of jobs that cannot be offshored. They are made up of the following sub-components: *iii)* face-to-face contact: face-to-face discussions, establishing and maintaining interpersonal relationships, assisting and caring for others, performing for or working directly with the public, coaching and developing others; *iv)* on-site job: inspecting equipment, structures, or material, handling and moving objects, controlling machines and processes, operating vehicles, mechanized devices, or equipment, and repairing and maintaining mechanical or electronic equipment; and *v)* decision-making: making decisions and solving problems, thinking creatively, developing objectives and strategies, responsibility for outcomes and results, frequency of decision making. With these refinements (over the DOT-DCP measure), we aim to capture the creative and strategic aspects of decision-making that can be less easily offshored.[20] We use the reverse of these measures of

Admittedly, they are higher between decision-making and information content (0.37) or automation (0.20).

[19] Appendix Table A2 lists the exact O*NET reference numbers of the generalized work activities and work context items that make up the five indexes and indicate the elements also used by JK and/or Blinder (2007).

[20] Admittedly, there are some non-offshorable service jobs that require little decision-making or little face-to-face interactions, but that need to be performed on-site. So components of our non-offshorability are best used together.

non-offshorability to capture "Offshorability".

For each occupation, the O*NET provides information on the "importance" and "level" of required work activity and on the frequency of five categorical levels of work context.[21] We follow Blinder (2009) in arbitrarily assigning a Cobb-Douglas weight of two thirds to "importance" and one third to "level" in a weighted sum for work activities. For "work contexts" elements, such as "frequency of decision making", we simply multiply the frequency by the value of the level. Each composite T_{jh} score for occupation j in category h is, thus, computed as

$$T_{jh} = \sum_{l=1}^{A_h} I_{jl}^{2/3} L_{jl}^{1/3} + \sum_{m=1}^{C_h} F_{jm} * V_{jm}, \tag{5}$$

where A_h is the number of work activity elements, and C_h the number of work context elements in the category T_{jh}, $h = 1, \ldots, 5$. Because the magnitude of the T_{jh} scores are not readily interpretable, we convert these scores to ordinal measures. We compute the quartiles of each of our five measures of task content and focus on an indicator variable for being in the upper quartile. For example, an occupation in the top 25% of our automation score will be classified as a routine job, to be compared to an occupation in the bottom 25%. Similarly, Autor and Dorn (2013) classify as "routine intensive" occupations those that fall in the top third of their RTI measure. With respect to our offshorability measure, our classification is consistent with Blinder (2009) whose own best guess is that 26% to 29% of U.S. jobs are potentially offshorable.[22]

Table 1 shows the percentage of workers, by gender, in five major occupational groups that rank in the top quartile of each of the five task content measures in 2000-02. These numbers are generally consistent with the evidence reported in related studies and show substantial differences across genders. For example, among professional and technical workers, men are more likely than women (74% vs. 42%) to be in the top quartile of our information content measure. Men are almost twice as likely as women (79% vs. 40%) to work off site when engaged in clerical and sales occupations. Women working in the primary sector, construction, and transportation sectors are in occupations more likely automated (69% vs. 37%) than men.

The highest percentage of workers in the top quartile of our information content measure are found among professional, managerial and technical occupations. More than three quarters of

[21]For example, the work context element "frequency of decision-making" has five categories: 1) never, 2) once a year or more but not every month, 3) once a month or more but not every week, 4) once a week or more but not every day, and 5) every day.

[22]We have experimented with continuous measures and other cut-offs and have not found that this affects our main findings.

production workers are in the top quartile of our automation/routine measure. Our two measures of technology thus separate well the effects of ICT innovation from plant floor automation. Consistent with expectations, none of our primary, construction, transport, or production workers, and very few service workers are in occupations ranked in the top quartile of the off-site measure. As a group, clerical and sales occupations are the ones most likely found in the top quartile of all three of task content measures linked to offshorability.[23]

As we noted above, there is a lack of consensus in the literature on how to group and interpret the effect of different task measures. In light of this, we present results based on some alternative groupings of task measures in addition to our main offshorability and technological change groupings discussed above. In Appendix Figure A6, we present an alternative decomposition where the information content of jobs (task *i)* above) and decision making (task *v)* above) are grouped under an "analytic and managerial content" category that are non-routine because they require creativity, professional judgment, etc. This category stands in contrast to the "automation/routine" task (a category on its own), and a modified offshorability grouping of tasks that includes tasks *iii)* (face-to-face contact) and *iv)* (on-site job). We also present separate results for each of our five basic task content measures in Appendix Tables A4 and A5, to allow readers to make their own evaluation.

III. Decomposing Changes in Distributions Using RIF-Regressions

In this section, we show how to formally decompose changes in the distribution of wages into the contribution of occupational tasks and other factors using the recentered influence function (RIF) regression approach introduced by Firpo, Fortin, and Lemieux (2009). As is well known, a standard regression can be used to perform a Oaxaca-Blinder decomposition for the mean of a distribution. RIF-regressions allow us to perform the same kind of decomposition for any distributional parameter, including percentiles.[24]

In general, any distributional parameter can be written as a functional $\nu(F_Y)$ of the cumulative distribution of wages, $F_Y(Y)$.[25] Examples include wage percentiles, the variance of log wage, the Gini coefficient, etc. The first part of the decomposition consists of dividing the overall change

[23]Including no decision-making in our offshorability measure thus allows a sharper distinction between managerial and clerical jobs.

[24]Firpo, Fortin, and Lemieux (2007) and Fortin, Lemieux and Firpo (2011) explain in more detail how to perform these decompositions, and show how to compute the standard errors for each element of the distribution. Here, we simply present a short summary of the methodology.

[25]In this section, we denote the wage using Y instead of W to be consistent with Firpo, Fortin, and Lemieux (2007) and the program evaluation literature.

in a given distributional parameter into a composition effect linked to changes in the distribution of the covariates, X, and a wage structure effect that reflects how the conditional distribution of wage $F(Y|X)$ changes over time. In a standard Oaxaca-Blinder decomposition, the wage structure effect only depends on changes in the conditional mean of wages, $\mathbb{E}(Y|X)$. More generally, however, the wage structure effect depends on the whole conditional wage distribution.

It is helpful to discuss the decomposition problem using the potential outcomes framework. We focus on differences in the wage distributions for two time periods, 1 and 0. For a worker i, let Y_{1i} be the wage that would be paid in period 1, and Y_{0i} the wage that would be paid in period 0. Therefore, for each i we can define the observed wage, Y_i, as $Y_i = Y_{1i} \cdot T_i + Y_{0i} \cdot (1 - T_i)$, where $T_i = 1$ if individual i is observed in period 1, and $T_i = 0$ if individual i is observed in period 0.

We use the notation $F_{Y_t|T=s}$ to indicate the distribution of wages that would prevail among workers observed in period s if they were paid under the wage structure of period t. For instance, $F_{Y_0|T=0}$ denotes the actual distribution in period 0, while $F_{Y_0|T=1}$ represents the counterfactual distribution that would have prevailed if workers in period 1 had been paid under the wage structure of period 0. There is also a vector of covariates $X \in \mathcal{X} \subset \mathbb{R}^K$ observed in both periods.

Consider Δ_O^ν, the overall change over time in the distributional statistic ν. We have

$$\Delta_O^\nu = \nu\left(F_{Y_1|T=1}\right) - \nu\left(F_{Y_0|T=0}\right) \tag{6}$$

$$= \underbrace{\nu\left(F_{Y_1|T=1}\right) - \nu\left(F_{Y_0|T=1}\right)}_{\Delta_S^\nu} + \underbrace{\nu\left(F_{Y_0|T=1}\right) - \nu\left(F_{Y_0|T=0}\right)}_{\Delta_X^\nu}, \tag{7}$$

where Δ_S^ν is the wage structure effect, while Δ_X^ν is the composition effect. Key to this decomposition is the counterfactual distributional statistics $\nu\left(F_{Y_0|T=1}\right)$. It represents the distributional statistic that would have prevailed if workers observed in the end period ($T = 1$) had been paid under the wage structure of period 0.

Estimating this type of counterfactual distribution is a well-known problem. For instance, DiNardo, Fortin and Lemieux (1996) suggest estimating this counterfactual by reweighting the period 0 data to have the same distribution of covariates as in period 1. We follow the same approach here, since Firpo, Fortin and Lemieux (2007) show that reweighting provides a consistent nonparametric estimate of the counterfactual distribution under the ignorability assumption.

However, the main goal of this paper is to separate the contribution of different subsets of covariates to Δ_O^ν , Δ_S^ν , and Δ_X^ν. This is easily done in the case of the mean where each component of the above decomposition can be written in terms of the regression coefficients and the mean of the covariates. For distributional statistics besides the mean, Firpo, Fortin, and Lemieux (2009) suggest estimating a similar regression where the usual outcome variable, Y, is replaced by the recentered influence function $\mathrm{RIF}(y;\nu)$ of the statistic ν. The recentering consists of adding back the distributional statistic ν to the influence function $\mathrm{IF}(y;\nu)$: $\mathrm{RIF}(y;\nu) = \nu + \mathrm{IF}(y;\nu)$, where the influence function is mean-zero by construction. Note that in the case of the mean where the influence function is $\mathrm{IF}(y;\mu) = y - \mu$, we have $\mathrm{RIF}(y;\mu) = \mu + (y - \mu) = y$. Since the RIF is the outcome variable y, the RIF-regression for the mean is a standard wage regression.

It is also possible to compute the influence function for many other distributional statistics. Of particular interest is the case of quantiles. The τ-th quantile of the distribution F is defined as the functional, $Q(F,\tau) = \inf\{y|F(y) \geq \tau\}$, or as q_τ for short. Its influence function is:

$$\mathrm{IF}(y;q_\tau) = \frac{\tau - \mathbb{1}\{y \leq q_\tau\}}{f_Y(q_\tau)}.$$

The recentered influence function of the τ^{th} quantile is $\mathrm{RIF}(y;q_\tau) = q_\tau + \mathrm{IF}(y;q_\tau)$.

Because of the law of iterated expectations, distributional statistics can be expressed in terms of expectations of the conditional recentered influence functions,

$$\nu(F) = \mathbb{E}_X\left[\mathbb{E}\left[\mathrm{RIF}(Y;\nu)|X\right]\right].$$

In particular, the τ^{th} quantile RIF-regression aggregates to the unconditional quantile of interest and captures both the between and within effects of the explanatory variables.

Now consider γ_t^ν, the vector coefficient from a linear projection of $\mathrm{RIF}(Y_t;\nu_t)$ on X given $T = t$

$$\gamma_t^\nu = \left(\mathbb{E}\left[XX^\intercal|T=t\right]\right)^{-1}\mathbb{E}\left[\mathrm{RIF}(Y_t;\nu_t)X|T=t\right], \quad t = 0, 1.$$

If the conditional expectation of RIF was linear, we could use the RIF-regression coefficients in an analog to the Oaxaca-Blinder decomposition. Firpo, Fortin and Lemieux (2007) point out, however, that there may be a bias in this type of decomposition if the non-linearity assumption does not hold. They propose a solution based on an hybrid approach that involves both reweighting and RIF-regressions, which also circumvents issues related to the dependence of RIF on ν. Their solution guarantees that, under ignorability, each decomposition term will only reflect either differences in wage structures or in the covariates distribution, but never both. Letting

$\nu_{01} = \nu\left(F_{Y_0|T=1}\right)$, we rewrite (7) as

$$\begin{aligned}
\Delta_S^{\nu} &= \mathbb{E}\left[\mathrm{RIF}(Y_1;\nu_1)|T=1\right] - \mathbb{E}\left[\mathrm{RIF}(Y_0;\nu_{01})|T=1\right] \\
\text{and} \quad \Delta_X^{\nu} &= \mathbb{E}\left[\mathrm{RIF}(Y_0;\nu_{01})|T=1\right] - \mathbb{E}\left[\mathrm{RIF}(Y_0;\nu_0)|T=0\right]
\end{aligned}$$

Reweighting allows us to write

$$\mathbb{E}\left[\mathbb{E}\left[\mathrm{RIF}(Y;\nu_{01})|X,T=0\right]X|T=1\right] = \mathbb{E}\left[\Psi(X)\mathbb{E}\left[\mathrm{RIF}(Y;\nu_{01})|X,T=0\right]X|T=0\right],$$

where the reweighting factor is given by $\Psi(X)$

$$\Psi(X) = \frac{\Pr(T=1|X)/\Pr(T=1)}{\Pr(T=0|X)/\Pr(T=0)}.$$

Since under the ignorability assumption,

$$\begin{aligned}
\gamma_{01}^{\nu} &= \left(\mathbb{E}\left[XX^{\intercal}|T=1\right]\right)^{-1}\mathbb{E}\left[\mathbb{E}\left[\mathrm{RIF}(Y_0;\nu_{01})|X,T=1\right]X|T=1\right] \\
&= \left(\mathbb{E}\left[XX^{\intercal}|T=1\right]\right)^{-1}\mathbb{E}\left[\mathbb{E}\left[\mathrm{RIF}(Y;\nu_{01})|X,T=0\right]X|T=1\right],
\end{aligned}$$

we can obtain a "reweighted" wage structure effect $\Delta_{S,p}^{\nu}$,

$$\begin{aligned}
\Delta_{S,p}^{\nu} &= \mathbb{E}\left[X|T=1\right]^{\intercal}\left(\gamma_1^{\nu} - \gamma_{01}^{\nu}\right) \\
&= \mathbb{E}\left[X|T=1\right]^{\intercal}\left(\mathbb{E}\left[XX^{\intercal}|T=1\right]\right)^{-1} \\
&\quad \cdot \mathbb{E}\left[X\left(\mathbb{E}\left[\mathrm{RIF}(Y;\nu_1)|X,T=1\right] - \mathbb{E}\left[\mathrm{RIF}(Y;\nu_{01})|X,T=0\right]\right)|T=1\right],
\end{aligned}$$

which is not influenced by differences between the distribution of X given $T=1$ and $T=0$, and reflects, under the ignorability assumption, a true change in the wage structure. This result would not hold if we were using $\mathbb{E}\left[X|T=1\right]^{\intercal}\left(\gamma_1^{\nu} - \gamma_0^{\nu}\right)$ as, in that case, differences in moments of X between the two groups would also affect the decomposition term.

An intuition for this result is that, for a given distribution of X, the approximation could change when the distribution of X changes even if the wage structure remains the same. For example, if the true relationship between Y and a single X is convex, the linear regression coefficient will increase when we shift the distribution of X up, even if the true (convex) wage structure remains unchanged. This means that γ_1^{ν} and γ_0^{ν} may be different just because they are estimated for different distributions of X even if the wage structure remains unchanged over time.

Estimation of these components is quite simple and is obtained by least squares.[26] The estimate of the composition effect $\widehat{\Delta}^{\nu}_{X,R}$ can be divided into a pure composition effect $\widehat{\Delta}^{\nu}_{X,p}$ using the wage structure of period 0 and a component measuring the specification error, $\widehat{\Delta}^{\nu}_{X,e}$:

$$
\begin{aligned}
\widehat{\Delta}^{\nu}_{X,R} &= \left(\overline{X}_{01} - \overline{X}_{0}\right)\widehat{\gamma}^{\nu}_{0} + \overline{X}_{01}\left[\widehat{\gamma}^{\nu}_{01} - \widehat{\gamma}^{\nu}_{0}\right] . \\
&= \qquad \widehat{\Delta}^{\nu}_{X,p} \qquad + \qquad \widehat{\Delta}^{\nu}_{X,e}
\end{aligned}
\tag{8}
$$

The specification error $\widehat{\Delta}^{\nu}_{X,e}$ captures the difference between $\widehat{\Delta}^{\nu}_{X,R}$, the composition effect estimated using a non-parametric reweighting approach (as in DiNardo, Fortin, and Lemieux, 1996), and the linear approximation $\widehat{\Delta}^{\nu}_{X,p}$ obtained using the RIF-regressions. A small specification error indicates that the RIF-regressions approximate well the "true" composition effect.

Similarly, the estimator of the wage structure effect can be written as

$$
\begin{aligned}
\widehat{\Delta}^{\nu}_{S,R} &= \overline{X}_{1}\left(\widehat{\gamma}^{\nu}_{1} - \widehat{\gamma}^{\nu}_{01}\right) + \left(\overline{X}_{1} - \overline{X}_{01}\right)\widehat{\gamma}^{\nu}_{01} \\
&= \qquad \widehat{\Delta}^{\nu}_{S,p} \qquad + \qquad \widehat{\Delta}^{\nu}_{S,e}
\end{aligned}
\tag{9}
$$

and reduces to the first term $\widehat{\Delta}^{\nu}_{S,p}$ as the reweighting error $\widehat{\Delta}^{\nu}_{S,e}$ goes to zero.[27]

This decomposition is easy to compute as it corresponds to two standard Oaxaca-Blinder decompositions performed on the estimated recentered influence functions. The first compares time period 0 and the reweighted time period 0 that mimics time period 1 and yields the pure composition effect. The second compares the time period 1 and the reweighted time period 0 and yields the pure wage structure effect.

IV. Decomposition Results: Task Content Variables vs. Other Factors

We now use our decomposition approach to look at the contribution of several explanatory factors to changes in the wage distribution. These factors consist of education (six categories), potential experience (nine categories), union coverage, marital status, race, and the quartile dummies for the five measures of occupational tasks discussed earlier. The various sets of factors are all included in the wage setting equation (4) presented in Section I. Education and experience

[26]The reweighting function is computed as the ratio of the predicted probabilities obtained from a logit specification that includes a rich set of interaction between the explanatory variables.

[27]The difference between $\overline{X}_{01}$ and $\overline{X}_{1}$ (the reweighting error) goes to zero in probability, as long as the estimate of $\Psi(.)$ converges to the true weighting function.

are part of the skills set S_i, occupational tasks measures are the T_{jh} variables that interact with skills, and the other variables are the Z_{it} factors that affect wages through other channels. Though these variables have a complex effect on the wage distribution, we are able to recover their impact using our flexible decomposition approach (See Appendix B for more discussion of this point).

Based on the evidence reported in Table 1, we focus our discussion on men, but present some complementary results for women. The challenge with estimating the impact of occupational tasks for women is that an important part of the variation in the occupational task measures comes from production operators, primary, construction and transport occupations. Forty-one percent of men but only 10 percent of women are in these occupations. Table 1 shows that these occupations have very low scores for automation and not on-site, and very high scores for automation and no face-to-face.

Before showing the decomposition results, it is useful to discuss some features of the estimated RIF-regression coefficients across the wage percentiles. As illustrative examples, we present the RIF-regression coefficients at each percentile of the male and female wage distribution for each of the four time periods for post-graduate education (base group is some college) and the top quartile (base group is the bottom quartile) of the five task content measures in Appendix Figure A3 and A4.[28] We also present an example of detailed regression estimates for a larger set of factors in Appendix Table A3 for the 1988-90 and 2000-02 periods for men. Note, however, that separate RIF-regressions are estimated for each gender, time period, and reweighted sample in the decomposition results below.

The figures show monotonic, and increasingly convex increases in the returns to a post-graduate degree over time, with more pronounced increases from 1988-90 to 2000-02. For the five task content measures, we generally find non-monotonic coefficients across the percentiles of the wage distribution that are qualitatively different for men and women. For men, "information" and "no face-to-face" have an inverse U-shaped impact, whereas for women only "information" has this shape. Interestingly, over time the peak positive impact of "information" on wages gradually moves from the lower to the upper end of the wage distribution. For women, the coefficients of "automation" exhibit a largely monotone decreasing curve across the wage distribution, with relatively small changes over time. For men, the coefficients of "automation" begin with a similarly monotone decreasing shape across the wage distribution in the late 1970s, and then become more and more U-shaped over time. This is consistent with Autor, Levy and Murnane

[28]More detailed estimates are available from the output files on the on-line data appendix.

(2003) who show that male workers in the middle of the distribution are more likely to experience negative wage changes as the "routine" tasks they used to perform can now be executed by robots (see also the discussion of equation (B-7) in Appendix B). These effects are qualitatively different from the impact of information content which is increasingly positive over time in the upper-middle part of the wage distribution for both men and women.[29]

Among the three task content measures associated with offshorability, the effect of "no face-to-face" displays the more important changes over time. For women, some negative effects in the top quartile in the late 1970s become positive in the 2000s and 2010s. For men, some positive effects in the top quartile decreased from the 2000-02 to 2011-12. Over the same period, men saw substantial decreases in the effect of "not on-site" in the bottom part of the wage distribution, whereas women saw substantial increases. Changes over time in the impact of the "no decision-making" measures appear less important.

As is well known (e.g. Oaxaca and Ransom, 1999), the detailed wage structure part of the decomposition depends arbitrarily on the choice of the base group, which we choose to be as neutral as possible to wage structure changes. The base group used in the RIF-regression models consists of non-union, white, and married workers with some college, 15 to 19 years of potential experience, and the bottom quartile of each of the five task measures.[30] A richer specification with additional interaction terms is used to estimate the logit models used in the computation of the reweighting factor.[31] The reweighting approach performs well in the sense that the reweighted means of the covariates for the base period are very close to those for the end period.[32]

A. Overall Decomposition Results

The results of the aggregate decomposition, which separates composition effects from wage structure effects, are presented in Table 2 for both men and women for each of the three time periods. We report changes over time in the 90-10 log wage differential as a measure of overall inequality, and changes in the 50-10 and 90-50 log wage differential as measures of low-end and

[29]At the very top end of the wage distribution, we find precipitously declining coefficients of "information" in 2011-12, that we have been able to trace to declining wages in occupations such as securities, commodities, and financial services sales agents (4820) and aircraft pilots and flight engineers (9030), for example.

[30]We use "some college" as the base group as it represents the modal education group in the 1990s and 2000s. For the 1976-78 to 1988-90 period, we use high schol graduates as the base group as it was still the modal education group during that period.

[31]The logit specification also includes a full set of interaction between experience and education, union status and education, union status and experience, and education and occupation task measures.

[32]The reweighting error is the second term in equation (9). If the reweighting was replicating the means perfectly, we would have $\overline{X}_1 = \overline{X}_{01}$ and the reweighting error would be equal to zero.

top-end wage inequality, respectively.[33] The numbers are multiplied by 100 to represent log points increases between the beginning and end period; given our large sample sizes, most of the results are statistically significant even if not always economically important.

We also report the specification errors for each time periods. Recall that the specification error is the difference between composition effects estimated using the RIF regressions (second term in equation (8)) and those estimated non-parametrically using a reweighted procedure. The specification errors are small relative to overall inequality changes, indicating that the linear RIF regressions provides a good approximation relative to non-paramteric estimates.[34]

Panel A presents the results for the 1976-78 to 1988-90 when, as shown Figures 1a and 1b, wage inequality increased monotonically (17 log points increase in the 90-10 gap for men, 34 for women). The second and third rows of the panel indicate that most of the increase (from 57 to 83 percent) comes from changes in the wage structure. The share of composition effects, that we link to increases in education and experience below, is also substantial, accounting for 12 to 56 percent of the total increase.

Consistent with Figure 1 and Autor, Katz and Kearney (2006), Panel B shows that inequality increased at the top-end (90-50) but decreased at the bottom-end (50-10) of the distribution between 1988-90 and 2000-02. The magnitude of the decline in wage inequality at the bottom is similar for men and women (about 8 log points). It is entirely attributable to wage structure effects, which account for more than 100 percent of the decline since composition effects go in the opposite direction. The latter are smaller for men leading to a more U-shaped pattern for that group. The increase in top-end wage inequality, 9 log points for men and 7 log points for women, is also mostly attributable to wage structure effects which account for 60 percent and 96 percent of the changes for men and women, respectively.

Panel C shows that inequality kept increasing steadily between 2000-02 and 2011-12. The 90-10 gap increased by about 10 log points for both men and women. Increases in the 90-50 gap were almost as large as in other decades, and were mostly due to wage structure effects (99 and 80 percent of the total change for men and women, respectively). By contrast, after declining during the 1990s, the 50-10 gap remained relatively stable between 2000-02 and 2011-12.

[33]Results for other measures overall wage inequality such the variance of log wages and the Gini coefficient are available upon request.

[34]One exception is the specification error for the 50-10 gap for women in the 1980s which exceeds one log point. This is likely due to the fact that the RIF regressions cannot fully capture the large change in the female wage distribution linked to the decline of the real value of the minimum wage over this period (DiNardo, Fortin, and Lemieux (1996), Lee (1999), Autor, Manning, and Smith (2010)). In an earlier version of the paper, we had introduced some corrections for the minimum wage and found that, except for the very bottom of the distribution in the 1976-78 to 1988-90 period, these adjustments did not change the substantive findings discussed below.

Figure 1 indicates, however, that the 50-10 fails to capture some interesting changes at the bottom end of the distribution that are consistent with polarization. For instance, wage at the very bottom (5th percentile and below) tend to increase during this period, while wages between the 15th and 30th percentiles fall substantially, especially for men. So while polarization in the 1990s was characterized by a relative symmetric U-shaped curves, the more recent changes look more like a slanted J-shaped curve.

A graphical version of the aggregate decomposition of Table 2 is presented in Appendix Figure A5. The graphs show that, for men, composition effects after year 2000 are similar (small positive effects) over the entire wage distribution, and cannot account for much of the change in inequality. Most the observed changes since the early 2000s are, therefore, due to wage structure effects. For women, composition effects are quite flat at the top end of the distribution, but continue to have a monotone increasing effect at the bottom end where they account for about 30 percent of the increase in wage dispersion.

Looking at all three decades, we conclude that, consistent with Lemieux (2006b), composition effects play a substantial role in inequality growth until the early 2000s. But changes in wage structure generally play a more important role, especially in terms of explaining the polarization of wages after the late 1980s. According, most of the remaining analysis will focus on the detailed wage structure effects linked to changes in returns to skills and occupational tasks.

B. Detailed Decomposition Results

The next step of the decomposition uses RIF-regressions to estimate the contribution of each set of explanatory factors to the composition and wage structure effects. The detailed decomposition results are presented in terms of changes in the 90-10, 50-10, and 90-50 log differentials in Table 3. The results by gender and time periods are regrouped by columns, while Panels 1) and 2) display composition effects and wage structure effects, respectively. To conserve space, we report the effect of union status, education and experience, as well as the two groupings of task measures, offshorability and technological change.[35] Note that the composition effects for the task content measures arise only from occupational changes across the decades and not from within-occupation changes in task content.[36]

[35]The effect of each set of categorical variables is obtained by summing up the contribution of the relevant covariates. For example, the effect for "education" is the sum of the effect of each of the five education categories shown in Appendix Table A1. For the task measures, we sum up (within the offshorability and technological change categories) the effect of being in the top quartile for each individual task measure (bottom quartile as base).

[36]To the extent that we are not able to fully ascertain the share of spurious factors behind the occupational changes, i.e. changes in occupation codes, changes in occupation labeling (from secretaries to administrative assistants), we refrain from

In the case of men, we report the results for all three subperiods. Although separating the effect of each occupational tasks measure is more challenging for women than men (see the above discussion), as supplementary evidence we show the results for women for the 1990s (1988-90 to 2000-02 period) when polarization was most pronounced.[37]

As the discussion below indicates, there are some limitations focusing on only three percentiles to understand distributional changes. More information about effects at each percentile is thus provided in Figures 3 and 4 for a reduced set of explanatory factors. The figures display smoothed changes to facilitate the visual interpretation of the results.[38]

Starting with composition effects for men, Panel 1 of Table 3 shows that unionization has consistent polarizing effects: it reduces wage inequality at the lower end and increases it at the upper end. However, the effect diminishes over time as the rate of unionization is falling at an increasingly smaller rate (from 30 percent in 1976-78 to 20 percent in 1988-90, 15 percent in 2000-02, and 12 percent in 2011-12). As a result, the contribution of de-unionization to inequality growth at the top end (90-50 gap) goes from 36 percent in the 1980s, to 29 percent in the 1990s, and 14 percent in the 2000s.[39] These effects are easier to see in Figure 3, where the composition effects linked to de-unionization are U-shaped and relatively more important than those linked to other factors. The fact that de-unionization plays an important role in inequality growth is consistent with earlier estimates (Freeman, 1993, Card, 1992, DiNardo, Fortin, and Lemieux, 1996). The new finding here is that de-unionization also contributes to the growth in labor market polarization. Consistent with Card, Lemieux, and Riddell (2004), de-unionization have no significant impacts on women.

Figure 3 also shows that the composition effects linked to education and experience have more uniform effects across the wage distribution. They play a less important role in male wage inequality growth than de-unionization, and have no impact on polarization. Composition effects

interpreting these composition effects. Luckily, these effects are generally small and often not significant.

[37]One limitation of the 2000-2002 to 2011-12 period is that there are two changes in occupational classification in 2003 and 2011. For men, results are qualitatively similar when we either use the whole 2000-12 to 2011-12 period or the more limited 2003-04 to 2009-10 period during which the occupational classification is stable. For women, however, results are sensitive to this choice, perhaps because of the problems we discuss earlier in the context of Table 1.

The same problem prevails (change in occupational classification) in the 1976-78 to 1988-90 period. In addition, results for women would also need to be interpreted with caution given the large increase in participation rates during this period (Mulligan and Rubinstein, 2008), and the overwhelming effect of the minimum wage. For these reasons, we only report here the detailed decomposition results for women in the 1988-90 to 2000-02 period. Results for the other periods are reported in Appendix Table A5.

[38]Note that the smoothing tends to overemphasize extreme tail effects (bottom or top 4 or 5 deciles). We discount these tail effects in our interpretation of the results.

[39]Note that, as in a Oaxaca-Blinder decomposition, these effects on the 90-50 (or other) gap can be computed directly by multiplying the percent decline in the unionization rate (Appendix Table A1) by the RIF-regression estimates of the union effects for 1988-90, for example (Appendix Table A4). Here they are obtained by dividing the numbers in the first row of Table 3 (columns 3, 6, and 9) by corresponding numbers for total changes in Table A2.

linked to education are more important for women, at least at the bottom end of the distribution where they account for a 3 log point growth in the 50-10 gap during the 1990s (similar results are obtained in the 1980s and 2000s). The large composition effects for women are not surprising in light of their large and ongoing increases in education. For instance, the percentage of female workers with a college (post-graduate) degree grew from 10(6) percent in 1978-78 to 14(8) percent in 1988-90, 20(9) percent in 2000-02, and 24(13) percent in 2011-12. The latest figures now exceed the level of education of men (21(11) percent with a college degree (post-graduate) degree).

The detailed wage structure effects are reported in Panel 2 of Table 3. The wage structure effects for education, experience, and our preferred grouping of the five task measures into a technology and an offshorability component are also reported in Figure 4. Considering workers' characteristics first, both Table 3 and Figure 4 show that changes in the wage structure linked to education play a substantial role at the top end of the distribution. In particular, over the 1990s, changes in the returns to education account for 60 percent of the growth in the 90-50 gap for men, and 81 percent for women. By contrast, changes in the returns to education play relatively minor roles at the bottom end of the distribution. These findings confirm Lemieux (2006a)'s conjecture that the large increase in the return to post-secondary education has contributed to a convexification of the wage distribution.

Changes in the wage structure linked to experience contribute to inequality growth in the 1980s, but are either not significant or go in the other direction in other periods, reflecting a progressive decline in the returns to experience since the mid-1980s that became even more dramatic in the 2000-02 to 2011-12 period. Union wage structure effects tend to be small except in the 1980s where they reduced wage dispersion at the top end of the distribution.

Turning to the role of occupational tasks, we begin by discussing the wage polarization of the 1990s which gave rise to the more nuanced view of technological change (Autor, Katz, and Kearney, 2008, Autor and Dorn, 2013), and called for a role for occupations. The results reported in Table 3 confirm that changes in the returns to occupational task measures linked to technological change (automation and information content) did play an important role in polarization. The estimated effect of -1.18 log points accounts for 21 percent of the decline in the 50-10 gap, while the effect at the top end (3.66 log points) accounts for 40 percent of the increase in the 90-50 gap. Thus, wage structure effects linked to technological change are at least as important as the composition effects linked to de-unionization in accounting for the polarization of male wages in the 1990s.

The more detailed decomposition for each of the five tasks measures (Appendix Tables A4 and A5) shows that most of the estimated polarization can be linked, as expected, to the automation/routine task. Likewise, for women changes in the returns to occupational task measures linked to technological change account for 36 percent of the increase in the 90-50 gap and 5 percent of the decline in the 50-10 gap during the 1990s. Appendix Table A5 shows that, for women, information content is playing a more important role in polarization than automation.

Figure 4 shows that, for both men and women in the 1990s, the negative effects linked to technological change (the solid curve) are most pronounced between the 20th and 40th percentiles. Focusing only the 50-10 gap hides some important changes happening between these two (arbitrary) percentiles. Interestingly, a closer examination of Figures 1a (men) and 1b (women) show that raw wage changes decline from the very bottom of the distribution up to the 20-30th percentiles, and then remain fairly stable until about the 60th where they start increasing again. Thus, technological change plays a particularly important role explaining changes up the 20-30th percentiles, in addition to accounting for some of the growth in inequality at the top end.[40] Again, this underlines the need to consider the whole distribution rather only specific percentiles in our decomposition.

Table 3 shows that offshorability is an important element of wage polarization for men, but not for women. In particular, the effect of 2.60 log points for men accounts for 29 percent of the growth in the 90-50 gap. Appendix Table A4 shows that changes in the returns to both the "no face-to-face" and "not onsite" tasks play a comparable role in the growth of the 90-50 gap. Offshorability also contribute to some of the decline in the 50-10 gap, though this effect is not statistically significant. In the case of women, the contribution of "no face to face" is comparable to the one for men, but this is offset by an opposite effect of "not onsite".[41] As a result, offshorability (the sum of the three task measures) plays a modest role in changes in the distribution of wages of women.

Figure 4 also shows that, for men in the 1990s, the wage structure effects linked to potentially offshorable jobs (the "o" line), although smaller in magnitude, parallel the technology effects. For men, both technological change and offshorability follow a distinct U-shape that closely mirrors the shape of the overall change in the wage distribution (Figure 1a). For women, we do not

[40] Computing the 40-10 instead yields an estimate of −2.1 log points for wage structure effects linked to information, which accounts for 77 percent of the change (−2.7) in this log wage differential over the 1990s for women.

[41] One possible explanation for this opposite finding for men and women is that jobs with the highest "onsite" scores are mostly traditional male jobs like construction, truck drivers, etc. (see Table 1). This is consistent with our earlier discussion where we argued that task measures aimed at capturing offshorability are better suited for men than women.

detect as meaningful offshorability effects.

Turning to the 1980s (for men), Table 3 shows that both technology and offshorability played an important role in the growth in top-end inequality during that period. The effect of technology (offshorability) on the 90-50 gap is 4.97 (3.67) log points, which account for 48 (35) percent of the total change. Interestingly, Appendix Table A4 shows that, during the 1980s, information content played a more important role than automation, which is the opposite of what we found in the 1990s. This suggests that the 1980s are better characterized by a conventional skill-biased technological change story, while "routine-biased" technological change became more important during the 1990s.

Importantly, Panel C of Figure 4 shows that the effect of both technology and offshorability is fairly monotone in the 1980s, but U-shaped during the 1990s. This shows that changes in returns to occupational tasks play an important role in explaining why inequality grew at all points of the distribution during the 1980s, and why the wage distribution became more polarized during the 1990s. Once again, an important part of these interesting changes in shape of the distribution are missed by just focusing on summary measures such as the 90-50 and 50-10 gaps.

The final set of results reported in Figure 4 (Panel D) indicate that technological change is no longer a significant factor in the increase of male wage inequality after year 2000. Instead, technological change effects are fairly similar across the male wage distribution. A second substantive finding for the 2000-02 to 20011-12 period is that offshorability has become a more important explanation for changes in male wage inequality. Figure 4 shows the wage structure effects linked to offshorability (the sum of no face-to-face, not onsite and no decision making) are large and negative in the upper half of the wage distribution.[42] This is consistent with some technical no face-to-face jobs involving little strategic decision making having suffered some wage decreases. The fact that offshorability has relatively more substantial effects on wages in recent years is consistent with the view of many labor market observers who have stressed the importance of offshoring, as opposed to technological change, in recent changes in the U.S. labor market.[43] Admittedly, as shown in Appendix Tables A4 and A5, our no decision-making measure accounts for a large part of the decrease in the 50-10, also consistent with either the type of offshoring effects described by Grossman and Rossi-Hansberg (2008) or the increase demand for services jobs argued by Autor and Dorn (2013).[44] This reflects the fact that measuring offshorability

[42]These are difficult to see in the 90-50 of Table 3 because the effects at the 50th and 90th centile are similar, masking more substantial negative effects in between.

[43]See Blinder (2007) and the refererences therein.

[44]See also Ngai and Pissarides (2007) who provide a theoretical perspective on the rise of service employment in industri-

potential using task measures is challenging, though these measures substantially improves the explanatory power of economic models.

A number of interesting conclusions emerge from our detailed wage decompositions. First, for men composition effects linked to de-unionization accounts for 24 and 29 percent, respectively, of the change in inequality at the lower (50-10) and upper (90-50) end of the distribution during the 1990s. For women, the lion's share of composition effects arise from increases in education levels which account for 58 percent of the change in overall (90-10) inequality over that period, although it does not help explain wage polarization. Second, the changing wage structure effects linked to education and the occupational task measures (both technological change and offshorability) all help account for the changing wage distribution during the 1980s and 1990s. In the 1990s, the wage structure effects linked to technological change account for 40 (36) percent in the increase in top end (90-50) male (female) wage inequality. But importantly, they also help account for decreasing wage inequality at the bottom, thus capturing quite well the different changes in shape that previously required rethinking of the sources of skill-biased technological change. Furthermore, the pattern of results is consistent with the view that technological change was skill-biased during the 1980s, "routine-biased" during the 1990s, but no longer played much of a role in the years 2000. By contrast, our results show that offshorability should be part of the discussion of changes in the wage structure, and are worth investigating in more detail in future work.

Finally, we present results based on an alternative set of grouping of tasks in Appendix Figure A6. As discussed at the end of Section II, information content and decision making are grouped under a new "analytical/managerial content" category, offshorability now only includes "no face-to-face" and "no onsite", while "automation/routine" is a category on its own.[45] In some cases using these three groupings does not affect our main conclusions, while in other cases it does. For instance, for men in the 1990s, "routine/automation", "offshorability", and "analytical/managerial content" contribute to polarization, just like technological change and offshorability did under our original grouping.

For men in the 1980s results are a bit noisy, but all three measures contribute to the overall increase in inequality, just as our two original groupings did for that period. For women in the 1990s the new grouping shows that "analytical/managerial content" is the most important con-

alized economies.

[45]In this alternative specification we do not use the reverse of the decision making task measure, as we did when it was part of the broader offshorability category.

tributor, while "technological change" was the dominant factor with the main grouping, reflecting the fact that information content is the dominant factor in that case. Men in the 2000-02 to 2011-12 period are arguably the case where this alternative grouping makes the biggest difference. The "analytical/managerial content" category has a mixed effect on inequality (positive at the bottom, negative at the top) and also reduces the contribution of offshorability to changes in the wage distribution.

Given the lack of consensus in the literature on how to group and interpret the effect of tasks, it is perhaps not surprising to find that the interpretation of some of the results changes when alternative groupings are used. But these interpretation issues do not affect our main conclusion that changes in returns to tasks have been a major contributor to changes in the wage distribution over the last few decades, and that they play a particularly important role in the polarization phenomenon of the 1990s.

V. Conclusion

In this paper, we show how changes in the return to occupational tasks contributed to changes in the distribution of wages. We present a wage setting model clarifying the connection between skills, tasks, and wages, and model changes in task prices as a function of detailed task content variables obtained using data from the O*NET. These task content variables are computed for each 3-digit occupation, and capture the extent to which occupations are potentially exposed to technological change (e.g. by being a "routine" occupation) or offshorability (e.g. by being an occupation with no "face-to-face" interactions with consumers). We use a decomposition procedure based on the influence function regression approach of Firpo, Fortin, and Lemieux (2009) to assess the role of changes in task prices and other factors in changes in wage inequality. The results indicate that changes in the return to task measures capturing offshorability and technological change played an important role in changes in the distribution of wages over the last three decades. In particular, they help explain the polarization of wages observed since the late 1990s for both men and women.

More generally, our results suggest that occupations and the task content of work play an important role in wage setting even after controlling for standard skill measures such as education. Like Acemoglu and Autor (2011), we conclude that it is essential to take account of tasks and occupations in our standard models of the labor market to adequately understand why the wage distribution has changed so much over the last few decades.

REFERENCES

Acemoglu, Daron and David H. Autor. 2011). "Skills, Tasks, and Technologies: Implications for Employment and Earnings." in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics,* Amsterdam: North-Holland, Vol. IV.B:. 1043-1172

Autor, David H. 2013. "The "Task Approach" to Labor Markets: An Overview." NBER Working Paper No. 18711, January

Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103(5): 1553-97

Autor, David H., Katz, Lawrence F. and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics,* 90(2): 300–323

Autor, David H., Katz, Lawrence F. and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96(2): 189–194

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content Of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics*, 118(4): 1279-1333

Autor, David H., Alan Manning and Christopher L. Smith. 2010. "The Contribution of the Minimum Wage to U.S. Wage Inequality over Three Decades: A Reassessment." NBER Working Paper No. 16533, November

Berman, Eli., John Bound, and Zvi Griliches. 1994. "Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers." *Quarterly Journal of Economics* 109(2): 367-397

Blinder, A. (2009) "How Many U.S. Jobs Might Be Offshorable?" *World Economics*, 10(2): 41 - 78

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

Card, David. 1992. "The Effects of Unions on the Distribution of Wages: Redistribution or Relabelling?" NBER Working Paper 4195

Card, David, T. Lemieux, and W. Craig Riddell. 2004. "Unions and Wage Inequality." *Journal of Labor Research* 25(4): 519-62

Caroli, Eve, and John Van Reenen. 2001. "Skill-Biased Organizational Change? Evidence From A Panel of British And French Establishments." *Quarterly Journal of Economics* 116(4): 1449-1492

Cortes, Guido Matias. 2012. "Where Have the Routine Workers Gone? A Study of Polarization Using Panel Data." University of Manchester, mimeo

Crinò, Rosario. 2010. "Service Offshoring and White-Collar Employment." *Review of Economic Studies,* 77(2): 595-632

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach." *Econometrica*, 64(6): 1001-1044

Ebenstein, Avraham, Ann Harrison, Margaret McMillan, and Shannon Phillips. 2013. "Estimating the Impact of Trade and Offshoring on American Workers Using the Current Population Surveys." forthcoming in the *Review of Economics and Statistics*

Feenstra, Robert, and Gordon Hanson. 2003. "Global Production Sharing and Inequality: A Survey of Trade and Wages." In E.K. Choi and J. Harrigan (eds) *Handbook of International Trade*, Oxford: Blackwell, 146-185

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. 2007. "Decomposing Distribution Using Recentered Influence Function Regressions." Unpublished manuscript, PUC-Rio and UBC

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. 2009. "Unconditional Quantile Regressions" *Econometrica* 77(3): 953-973

Fortin, Nicole M., Thomas Lemieux, and Sergio Firpo. 2011. "Decomposition Methods in Economics." in O. Ashenfelter and D. Card, eds., *Handbook of Economics,* Amsterdam: North-Holland, Vol. IV.A: 1-102

Freeman, Richard B. 1993. "How Much has Deunionization Contributed to the Rise of Male Earnings Inequality?" In Sheldon Danziger and Peter Gottschalk, eds. *Uneven Tides: Rising Income Inequality in America.* New York: Russell Sage Foundation, 133-63

Freeman, Richard B. 1995. "Are Your Wages Set in Beijing?" *Journal of Economic Perspectives*, 9(3): 15-32

Gibbons, R., L.F. Katz, T. Lemieux, and D. Parent. 2005. "Comparative Advantage, Learning, and Sectoral Wage Determination." *Journal of Labor Economics* 23(4): 681-724

Goos, Maarten, and Alan Manning. 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *Review of Economics and Statistics*, 89(1): 118-133

Goos, Maarten, Alan Manning and Anna Salomons. 2009. "The Polarization of the European Labor Market." *American Economic Review,* 99(2): 58-63

Goos, Maarten, Alan Manning, and Anna Salomons. 2010. "Recent Changes in the European Employment Structure: The Roles of Technological Change, Globalization and Institutions." unpublished manuscript

Grossman, Gene M., and Esteban Rossi-Hansberg. 2008. "Trading Tasks: A Simple Theory of Offshoring." *American Economic Review,* 98(5): 1978—97

Heckman, James J., and Jose Scheinkman. 1987. "The Importance of Bundling in a Gorman-Lancaster Model of Earnings." *Review of Economic Studies*, 54(2): 243-55

Jensen, J. Bradford and Lori G. Kletzer. 2010. "Measuring the Task Content of Offshorable Services Jobs, Tradable Services and Job Loss." In *Labor in the New Economy*, Katharine Abraham, Mike Harper, and James Spletzer, eds., Chicago: University of Chicago Press. 309-335

Jung, Jaewon and Jean Mercenier. 2010. "Routinization-Biased Technical Change, Globalization and Labor Market Polarization: Does Theory Fit the Facts?" ERMES Working paper, Université Panthéon-Assas (Paris 2)

Katz, Lawrence F. and Kevin M. Murphy (1992) "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *Quarterly Journal of Economics* 117(5): 914-940

Lazear, Edward P. 2009. "Firm-Specific Human Capital: A Skill-Weights Approach." *Journal of Political Economy* 117(5): 914-940

Lee, David S. 1999. "Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage." *Quarterly Journal of Economics* 114(3): 977-1023

LEMIEUX, THOMAS. 2006a. "Post-secondary Education and Increasing Wage Inequality." *American Economic Review* 96(2): 195-199

LEMIEUX, THOMAS. 2006b. "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *American Economic Review* 96(3): 461-498

MICHAELS, GUY, ASHWINI NATRAJ, AND JOHN VAN REENEN. 2013. "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over 25 years." forthcoming in the *Review of Economics and Statistics*

MULLIGAN, CASEY B. AND YONA RUBINSTEIN. 2008. "Selection, Investment, and Women's Relative Wages Over Time." *Quarterly Journal of Economics* 123(3): 1061-1110

MURPHY, K.M., AND F. WELCH. 1991. "The Role of International Trade in Wage Differentials." In M. Kosters (ed.) *Workers and Their Wages*, Washington DC: American Enterprise Institute Press, 39-69

NGAI, L. RACHEL, AND CHRISTOPHER A. PISSARIDES. 2007. "Structural change in a multi-sector model of growth." *American Economic Review* 97(1): 429-443

OAXACA, R. AND M.R. RANSOM. 1999. "Identification in Detailed Wage Decompositions." *Review of Economics and Statistics* 81(1): 154–157

POLETAEV, MAXIM AND CHRIS ROBINSON. 2008. "Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984-2000." *Journal of Labor Economics* 26(3): 387-420

ROSEN, SHERWIN. 1978. "Substitution and Division of Labour." *Economica* 45(179): 235-50.

YAMAGUCHI, SHINTARO. 2012. "Tasks and Heterogeneous Human Capital." *Journal of Labor Economics*, 30(1): 1-53.

Table 1. Percentage of Workers in the Top Quartile of O*NET Indexes by Major Occupation Group in 2000/02

O*NET Indexes	Percentage of workers		Technology				Offshorability					
			Information		Automation		Not On-Site		No Face-to-Face		No Decision	
	Men	Women	Men	Women	Men	Women	Men	Women	Men	Women	Men	Women
Overall	100	100	28	26	27	27	25	25	26	25	28	27
Professional, Managerial, and Technical	32	38	74	42	4	4	37	24	19	13	0	0
Clerical, Sales	16	36	18	20	28	48	79	40	4	17	49	48
Production, Operators	11	6	1	0	82	80	0	0	79	91	42	37
Primary, Construction, and Transport	30	4	3	1	37	69	0	0	25	71	30	16
Service	11	16	12	20	3	7	5	9	27	37	63	46

Note: The numbers in each of the five O*NET indexes columns indicate the percentage of workers in each major occupation group by gender, which fall in the top 75 percent of their category.

Table 2. Aggregate Decomposition Results

	Men			Women		
Inequality Measure:	90-10	50-10	90-50	90-10	50-10	90-50
A: 1976/78 to 1988/90						
Total Change	17.42***	7.00***	10.43***	33.92***	24.30***	9.63***
	(0.92)	(0.8)	(0.56)	(0.46)	(0.48)	(0.45)
Wage Structure	11.81***	5.82***	5.99***	25.70***	18.42***	7.28***
	(1.25)	(1.14)	(1.07)	(0.86)	(0.9)	(0.91)
Composition	7.53***	1.72***	5.81***	6.76***	2.93***	3.82***
	(0.46)	(0.36)	(0.36)	(0.45)	(0.33)	(0.43)
Specification Error	-0.83	0.02	-0.84	2.43***	3.16***	-0.74
	(1.05)	(0.98)	(0.74)	(0.68)	(0.77)	(0.85)
B: 1988/90 to 2000/02						
Total Change	3.49***	-5.61***	9.11***	3.81***	-3.23***	7.04***
	(0.56)	(0.48)	(0.45)	(0.35)	(0.29)	(0.3)
Wage Structure	-2.90***	-8.36***	5.46***	-1.43***	-8.20***	6.78***
	(0.5)	(0.51)	(0.45)	(0.42)	(0.26)	(0.37)
Composition	4.78***	2.17***	2.61***	5.06***	4.55***	0.51***
	(0.21)	(0.14)	(0.18)	(0.21)	(0.16)	(0.16)
Specification Error	1.70***	0.67**	1.03***	0.18	0.44**	-0.26
	(0.26)	(0.26)	(0.2)	(0.21)	(0.17)	(0.22)
C: 2000/02 to 2011/12						
Total Change	8.87***	0.54	8.33***	9.51***	2.67***	6.84***
	(0.63)	(0.49)	(0.67)	(0.43)	(0.33)	(0.44)
Wage Structure	10.12***	1.84***	8.28***	6.10***	0.62*	5.48***
	(0.61)	(0.41)	(0.68)	(0.51)	(0.33)	(0.54)
Composition	-0.04	-0.53***	0.49**	3.73***	2.52***	1.21***
	(0.26)	(0.15)	(0.2)	(0.2)	(0.15)	(0.18)
Specification Error	-1.32***	-0.71***	-0.61**	-0.56**	-0.57***	0.01
	(0.27)	(0.25)	(0.3)	(0.27)	(0.18)	(0.25)

Note: Log wage differentials × 100. Bootstrapped standard errors are in parentheses (100 replications of the entire procedure). The formulas for the different components are the following and the difference between the total change and the sum of the three components shown is the reweighting error (not shown). Asterisks indicate statistical significance at the 1% (***), 5% (**) or 10%(*) level.

Total Change: $\hat{\Delta}_O^{\nu} = \overline{\mathrm{R\hat{I}F}(Y_1, \nu)} - \overline{\mathrm{R\hat{I}F}(Y_0, \nu)}$

Composition : $\hat{\Delta}_{X,p}^{\nu} = (\overline{X}_{01} - \overline{X}_0)\hat{\gamma}_0^{\nu}$

Wage Structure: $\hat{\Delta}_{S,p}^{\nu} = \overline{X}_1(\hat{\gamma}_1^{\nu} - \hat{\gamma}_{01}^{\nu})$

Specification Error: $\hat{\Delta}_{X,e}^{\nu} = \overline{X}_{01}\,(\hat{\gamma}_{01}^{\nu} - \hat{\gamma}_0^{\nu})$

Table 3. Detailed Decomposition Results - Main Factors

	Men									Women		
Years	A: 1976/78 to 1988/90			B: 1988/90 to 2000/02			C: 2000/02 to 2011/12			D: 1988/90 to 2000/02		
Inequality Measure:	90-10	50-10	90-50	90-10	50-10	90-50	90-10	50-10	90-50	90-10	50-10	90-50
1) Detailed Composition Effects:												
Union	2.03***	-1.73***	3.76***	1.29***	-1.35***	2.64***	0.53***	-0.67***	1.20***	-0.01	-0.13***	0.13***
	(0.11)	(0.1)	(0.13)	(0.05)	(0.05)	(0.08)	(0.04)	(0.04)	(0.07)	(0.01)	(0.26)	(0.03)
Education	0.81***	0.34***	0.46***	-0.18	0.81***	-1.00***	1.19***	0.69***	0.50***	2.93***	2.96***	-0.03
	(0.21)	(0.13)	(0.14)	(0.12)	(0.08)	(0.09)	(0.18)	(0.09)	(0.12)	(0.14)	(0.09)	(0.11)
Experience	-0.03	0.44***	-0.47***	1.51***	1.21***	0.30***	-0.10	0.06	-0.16*	0.29***	0.44***	-0.15**
	(0.14)	(0.1)	(0.1)	(0.1)	(0.06)	(0.09)	(0.09)	(0.07)	(0.06)	(0.06)	(0.04)	(0.06)
Technology	-1.93***	-0.31	-1.82***	-0.21***	0.00	-0.22***	0.10**	-0.39***	0.49***	0.79***	0.86***	-0.06
	(0.3)	(0.21)	(0.28)	(0.08)	(0.06)	(0.07)	(0.05)	(0.05)	(0.06)	(0.23)	(0.12)	(0.15)
Offshoring	-0.13	0.42**	-0.55***	0.04	0.24***	-0.19***	0.32***	0.32***	0.00	0.48***	0.24***	-0.26***
	(0.25)	(0.2)	(0.18)	(0.04)	(0.04)	(0.03)	(0.07)	(0.05)	(0.05)	(0.06)	(0.05)	(0.04)
2) Detailed Wage Structure Effects												
Union	-1.27***	0.80*	-2.07***	0.75***	0.24*	0.52***	0.35*	0.71***	-0.36*	-0.37**	-0.43***	0.07
	(0.35)	(0.41)	(0.39)	(0.16)	(0.14)	(0.17)	(0.18)	(0.15)	(0.22)	(0.18)	(0.1)	(0.19)
Education	4.98***	2.66**	2.32	6.07***	0.58	5.49***	2.28**	0.62	1.66*	7.38***	1.68***	5.70***
	(1.64)	(1.19)	(1.55)	(0.77)	(0.56)	(0.77)	(0.95)	(0.66)	(0.89)	(0.67)	(0.52)	(0.6)
Experience	3.79	1.81	1.98	-1.78	-0.67	-1.11	-4.08***	0.01	-4.09***	-1.52	-0.35	-1.17
	(2.48)	(1.75)	(2.08)	(1.28)	(0.85)	(1.16)	(1.45)	(1.02)	(1.46)	(1.18)	(0.75)	(1.04)
Technology	4.14***	-0.84	4.98***	2.48***	-1.18**	3.66**	-0.64	1.20**	-1.84***	2.37***	-0.16	2.53***
	(1.26)	(1.23)	(1.23)	(0.62)	(0.5)	(0.53)	(0.74)	(0.57)	(0.69)	(0.51)	(0.47)	(0.44)
Offshoring	0.19	1.31	-1.12	2.03***	-0.56	2.59***	-2.69**	-2.72***	0.03	0.35	0.12	0.22
	(1.36)	(1.09)	(1.12)	(0.66)	(0.52)	(0.57)	(1.07)	(0.61)	(0.96)	(0.58)	(0.38)	(0.52)

Note: Log wage differentials × 100. Bootstrapped standard errors are in parentheses (100 replications of the entire procedure). Explanatory variables include dummies for union coverage, married, non-white, 6 education classes, 9 experience classes, 4 quartiles of the task indexes. The reported effect for the task indexes is for the upper quartile when the bottom quartile is omitted. Asterisks indicate statistical significance at the 1% (***), 5% (**) or 10%(*) level.

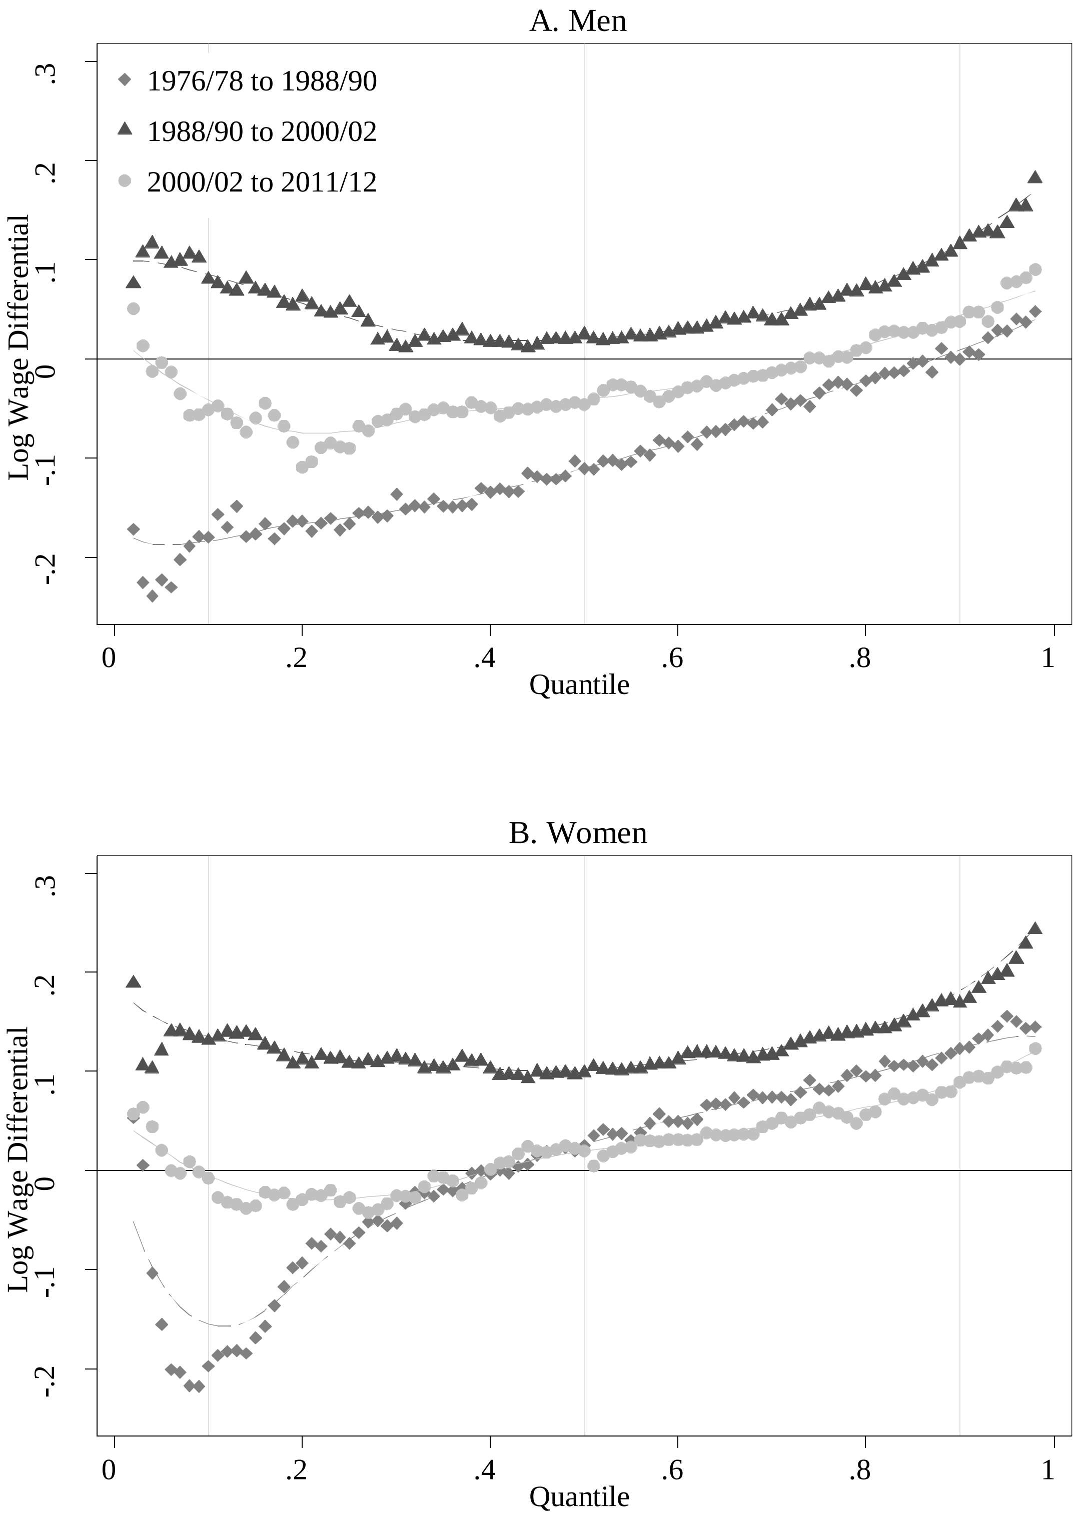

Figure 1. Changes in Real Log Wages by Percentile

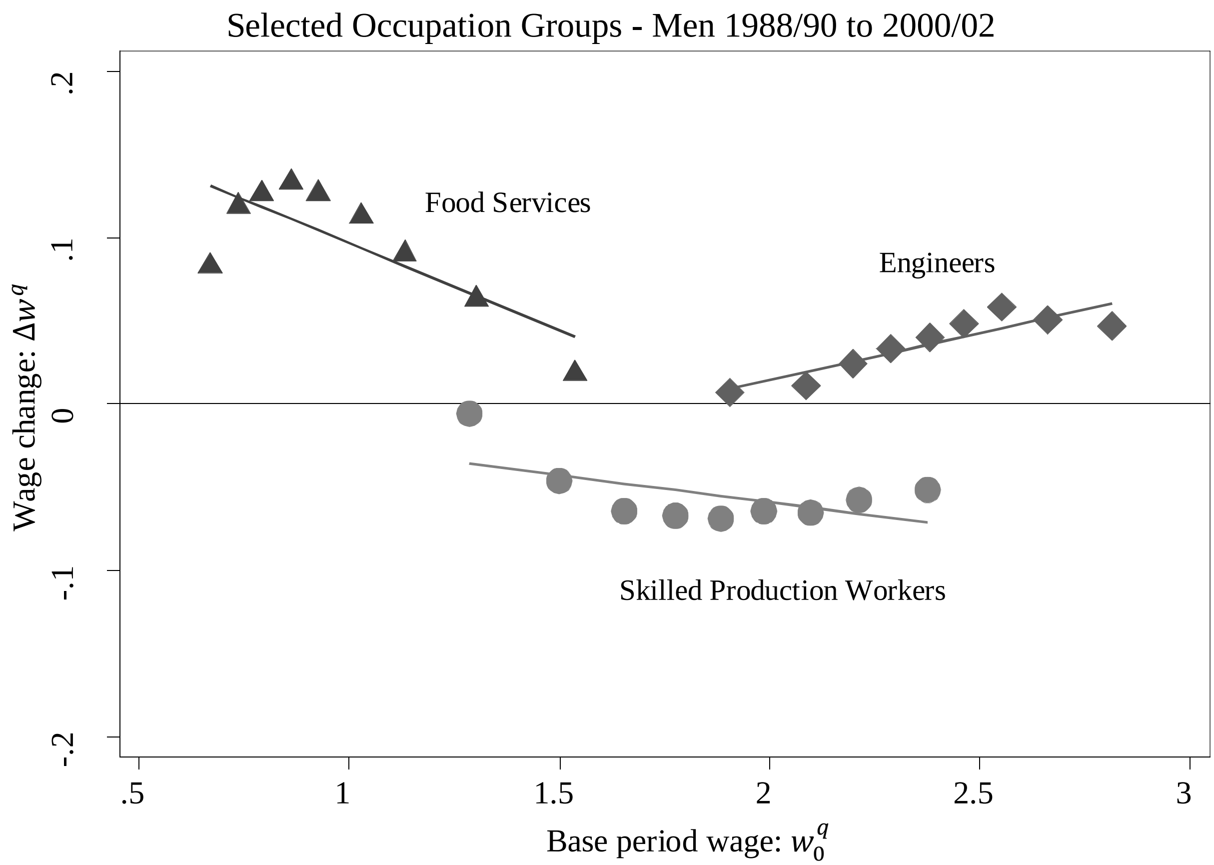

Figure 2. Occupation-Specific Wage Changes by Decile

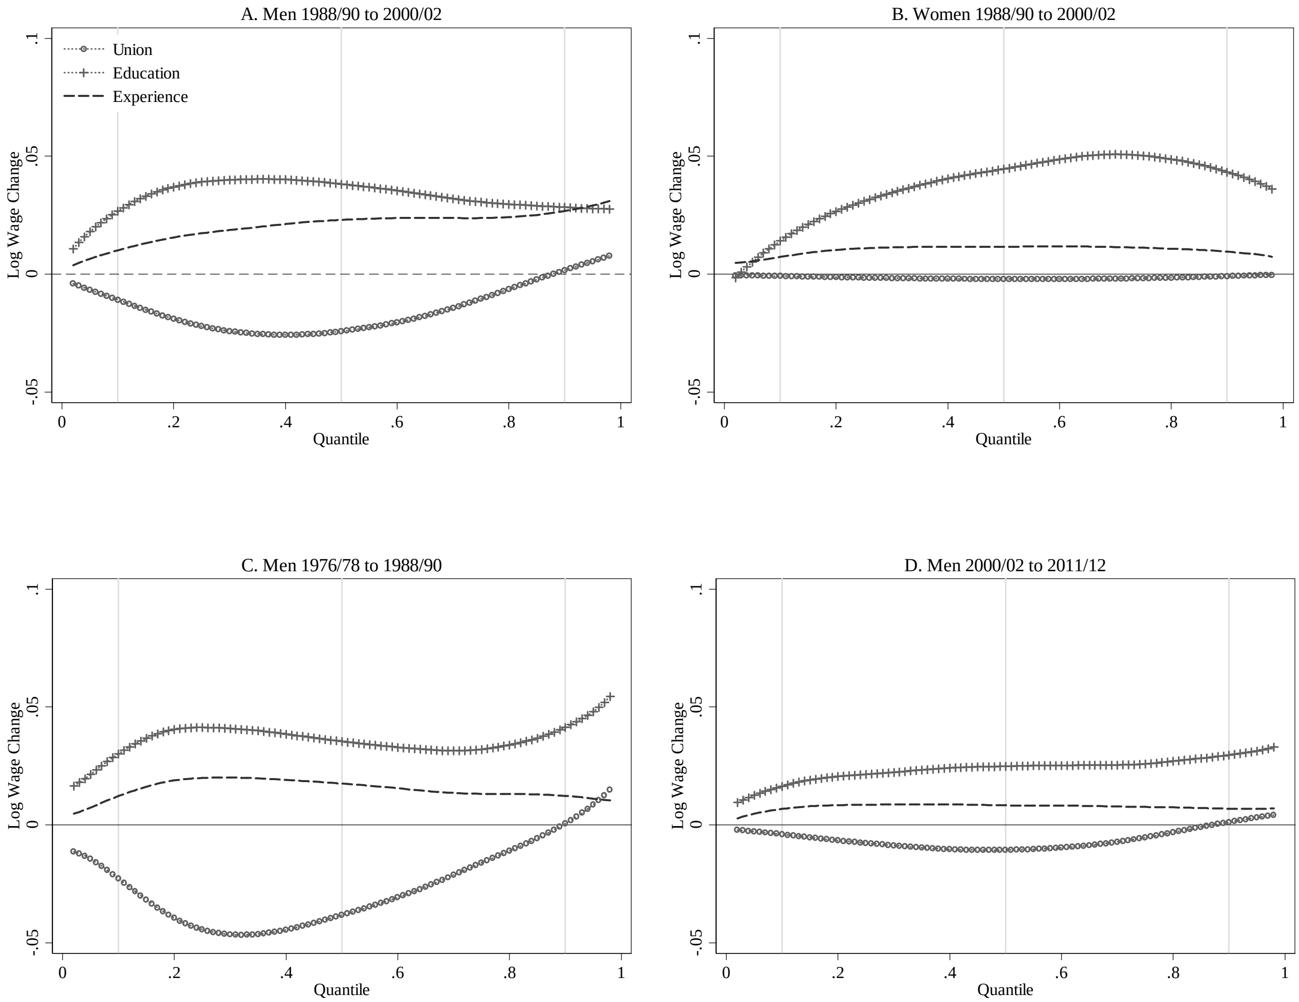

Figure 3. Detailed Decomposition of Composition Effects – Traditional Factors

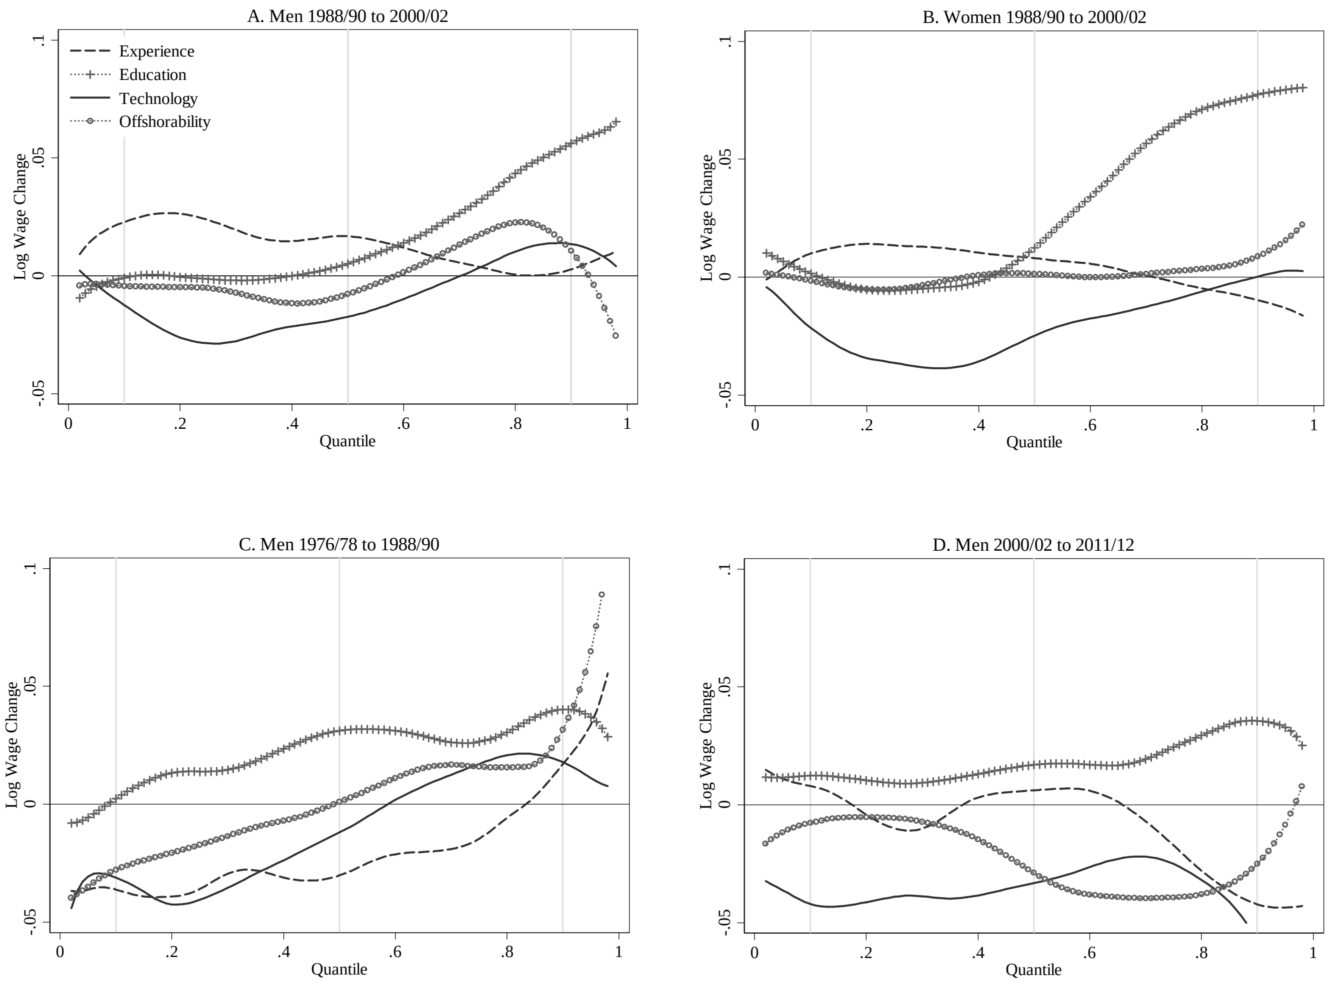

Figure 4. Detailed Decomposition of Wage Structure Effects – Selected Groupings

NOT INTENDED FOR PUBLICATION

APPENDIXES TO

Occupational Tasks and Changes in the Wage Structure

by Sergio Firpo, Nicole M. Fortin, and Thomas Lemieux

Appendix A - Supplementary Tables and Figures

Appendix Table A1. Descriptive Statistics

Men	1976/78		1988/90		Difference	2000/02		Difference	2011/12		Difference
	Means	Standard Deviation	Means	Standard Deviation	in Means (88/90–76/78)	Means	Standard Deviation	in Means (00/02-88/90)	Means	Standard Deviation	in Means (11/12–00/02)
Log wages	1.851	0.520	1.753	0.583	-0.098	1.812	0.597	0.059	1.787	0.623	-0.025
Union covered	0.295	0.456	0.202	0.401	-0.093	0.149	0.356	-0.053	0.123	0.329	-0.025
Non-white	0.101	0.302	0.127	0.333	0.026	0.140	0.347	0.013	0.154	0.361	0.014
Non-Married	0.295	0.456	0.386	0.487	0.091	0.415	0.493	0.028	0.439	0.496	0.024
Education											
Primary	0.103	0.303	0.060	0.237	-0.043	0.042	0.200	-0.018	0.036	0.187	-0.006
Some HS	0.174	0.379	0.121	0.326	-0.053	0.089	0.285	-0.032	0.062	0.241	-0.027
High School	0.369	0.483	0.379	0.485	0.009	0.312	0.463	-0.067	0.300	0.458	-0.011
Some College	0.168	0.374	0.203	0.402	0.035	0.274	0.446	0.071	0.279	0.449	0.005
College	0.106	0.307	0.137	0.344	0.032	0.188	0.391	0.051	0.212	0.409	0.024
Post-grad	0.080	0.272	0.100	0.301	0.020	0.095	0.294	-0.005	0.111	0.314	0.015
Age	35.708	12.854	35.766	11.738	0.058	37.569	11.824	1.803	39.367	12.453	1.798
No. of Observations	70516		226076			167929			100807		

Women	1976/78		1988/90		Difference	2000/02		Difference	2011/12		Difference
	Means	Standard Deviation	Means	Standard Deviation	in Means (88/90–76/78)	Means	Standard Deviation	in Means (00/02–88/90)	Means	Standard Deviation	in Means (11/12–00/02)
Log wages	1.463	0.444	1.471	0.522	0.008	1.602	0.547	0.131	1.628	0.581	0.025
Union covered	0.153	0.360	0.129	0.335	-0.024	0.122	0.327	-0.007	0.114	0.318	-0.007
Non-white	0.127	0.333	0.145	0.352	0.018	0.165	0.372	0.020	0.179	0.384	0.014
Non-Married	0.429	0.495	0.451	0.498	0.022	0.472	0.499	0.021	0.488	0.500	0.016
Education											
Primary	0.069	0.254	0.034	0.180	-0.036	0.024	0.154	-0.009	0.020	0.141	-0.004
Some HS	0.163	0.369	0.101	0.301	-0.062	0.073	0.261	-0.027	0.047	0.211	-0.027
High School	0.450	0.498	0.418	0.493	-0.033	0.303	0.459	-0.115	0.247	0.431	-0.056
Some College	0.165	0.371	0.232	0.422	0.067	0.315	0.465	0.084	0.322	0.467	0.007
College	0.097	0.296	0.135	0.342	0.038	0.196	0.397	0.061	0.236	0.425	0.040
Post-grad	0.055	0.228	0.082	0.274	0.027	0.089	0.285	0.007	0.128	0.334	0.039
Age	35.135	13.053	35.894	11.838	0.759	37.965	12.019	2.071	39.946	12.718	1.981
No. of Observations	54246		212489			166154			100326		

Appendix Table A2. O*NET Work Activities & Work Context

A) Characteristics linked to Technological Change/Offshorability	
1) Information Content	
4.A.1.a.1	Getting Informations (JK)
4.A.2.a.2	Processing Informations (JK)
4.A.2.a.4	Analyzing Data or Informations (JK)
4.A.3.b.1	Interacting With Computers (JK)
4.A.3.b.6	Documenting/Recording Informations (JK)
2) Automation/Routine	
4.C.3.b.2	Degree of Automation
4.C.3.b.7	Importance of Repeating Same Tasks
4.C.3.b.8	Structured versus Unstructured Work (reverse)
4.C.3.d.3	Pace Determined by Speed of Equipment
4.C.2.d.1.i	Spend Time Making Repetitive Motions
B) Characteristics linked to Non-Offshorability	
3) Face-to-Face Contact	
4.C.1.a.2.l	Face-to-Face Discussions
4.A.4.a.4	Establishing and Maintaining Interpersonal Relationships (JK,B)
4.A.4.a.5	Assisting and Caring for Others (JK,B)
4.A.4.a.8	Performing for or Working Directly with the Public (JK,B)
4.A.4.b.5	Coaching and Developing Others (B)
4) On-site Job	
4.A.1.b.2	Inspecting Equipment, Structures, or Material (JK)
4.A.3.a.2	Handling and Moving Objects
4.A.3.a.3	Controlling Machines and Processes
4.A.3.a.4	Operating Vehicles, Mechanized Devices, or Equipment
4.A.3.b.4	Repairing and Maintaining Mechanical Equipment (*0.5)
4.A.3.b.5	Repairing and Maintaining Electronic Equipment (*0.5)
5) Decision-Making	
4.A.2.b.1	Making Decisions and Solving Problems (JK)
4.A.2.b.2	Thinking Creatively (JK)
4.A.2.b.4	Developing Objectives and Strategies
4.C.1.c.2	Responsibility for Outcomes and Results
4.C.3.a.2.b	Frequency of Decision Making

Note: (JK) indicates a work activity used in Jensen and Kletzer (2007), (B) a work activity used or suggested in Blinder (2007).

Appendix Table A3. RIF-Regression Coefficients on Male Log Wages

Explanatory Variables	Years: Quantiles:	1988/90 10	1988/90 50	1988/90 90	2000/02 10	2000/02 50	2000/02 90
Union covered		0.218	0.454	-0.048	0.161	0.414	-0.091
		(0.003)	(0.005)	(0.005)	(0.004)	(0.006)	(0.007)
Non-white		-0.070	-0.136	-0.080	-0.037	-0.126	-0.045
		(0.008)	(0.005)	(0.005)	(0.006)	(0.005)	(0.008)
Non-Married		-0.152	-0.127	-0.036	-0.095	-0.142	-0.089
		(0.005)	(0.004)	(0.005)	(0.004)	(0.005)	(0.006)
Education (Some College omitted)							
Primary		-0.443	-0.504	-0.220	-0.496	-0.519	-0.134
		(0.014)	(0.008)	(0.006)	(0.018)	(0.01)	(0.007)
Some HS		-0.431	-0.271	-0.089	-0.443	-0.300	-0.015
		(0.011)	(0.006)	(0.005)	(0.01)	(0.006)	(0.005)
High School		-0.051	-0.134	-0.106	-0.047	-0.157	-0.072
		(0.006)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)
College		0.103	0.220	0.338	0.063	0.248	0.449
		(0.005)	(0.006)	(0.009)	(0.005)	(0.006)	(0.01)
Post-grad		0.042	0.230	0.665	0.022	0.278	1.006
		(0.005)	(0.006)	(0.014)	(0.006)	(0.007)	(0.022)
Potential Experience (15< Experience < 20 omitted)							
Experience <5		-0.559	-0.472	-0.337	-0.438	-0.414	-0.247
		(0.009)	(0.007)	(0.009)	(0.009)	(0.008)	(0.01)
5< Experience < 10		-0.067	-0.283	-0.303	-0.062	-0.259	-0.278
		(0.007)	(0.006)	(0.009)	(0.007)	(0.007)	(0.011)
10< Experience < 15		-0.016	-0.127	-0.190	-0.027	-0.108	-0.140
		(0.007)	(0.006)	(0.01)	(0.007)	(0.007)	(0.013)
20< Experience < 25		-0.002	-0.054	-0.102	-0.012	-0.049	-0.030
		(0.006)	(0.006)	(0.01)	(0.006)	(0.007)	(0.011)
25< Experience < 30		(0.01)	(0.032)	(0.06)	-(0.002)	(0.023)	(0.007)
		(0.006)	(0.007)	(0.012)	(0.006)	(0.007)	(0.012)
30< Experience < 35		0.017	0.045	0.060	-0.003	0.023	0.019
		(0.008)	(0.008)	(0.011)	(0.006)	(0.007)	(0.014)
35< Experience < 40		0.023	0.021	0.048	0.004	0.008	0.035
		(0.009)	(0.008)	(0.012)	(0.007)	(0.009)	(0.015)
Experience > 40		0.085	0.015	-0.027	-0.007	-0.044	-0.042
		(0.009)	(0.008)	(0.01)	(0.012)	(0.01)	(0.014)
O*NET Measures							
Information Content		0.067	0.086	0.023	0.052	0.096	0.044
		(0.003)	(0.002)	(0.004)	(0.003)	(0.003)	(0.004)
Automation		0.015	-0.035	-0.044	0.014	-0.055	-0.023
		(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
No Face-to-Face		0.114	0.122	0.115	0.086	0.120	0.121
		(0.003)	(0.002)	(0.003)	(0.002)	(0.002)	(0.005)
Non On-Site Job		-0.027	0.050	0.092	-0.028	0.044	0.104
		(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)
No Decision-Making		-0.157	-0.148	-0.142	-0.136	-0.157	-0.137
		(0.004)	(0.003)	(0.004)	(0.003)	(0.004)	(0.005)
Constant		1.598	2.774	2.465	1.219	1.896	2.524
		(0.011)	(0.012)	(0.015)	(0.006)	(0.006)	(0.01)
Number of obs.			226,076			167,929	

Note: Bootstrapped standard errors are in parentheses (100 replications of the entire procedure).

Appendix Table A4. Detailed Decomposition Results - Main Factors: Men

Years	A: 1976/78 to 1988/90			B: 1988/90 to 2000/02			C: 2000/02 to 2011/12		
Inequality Measure:	90-10	50-10	90-50	90-10	50-10	90-50	90-10	50-10	90-50
1) Detailed Composition Effects:									
Union	2.03***	-1.73***	3.76***	1.29***	-1.35***	2.64***	0.53***	-0.67***	1.20***
	(0.11)	(0.1)	(0.13)	(0.05)	(0.05)	(0.08)	(0.04)	(0.04)	(0.07)
Education	0.81***	0.34***	0.46***	-0.18	0.81***	-1.00***	1.19***	0.69***	0.50***
	(0.21)	(0.13)	(0.14)	(0.12)	(0.08)	(0.09)	(0.18)	(0.09)	(0.12)
Experience	-0.03	0.44***	-0.47***	1.51***	1.21***	0.30***	-0.10	0.06	-0.16*
	(0.14)	(0.1)	(0.1)	(0.1)	(0.06)	(0.09)	(0.09)	(0.07)	(0.06)
Info	-2.43***	-0.31	-2.12***	-0.37***	0.02	-0.39***	0.11***	-0.26***	0.37***
	(0.26)	(0.2)	(0.26)	(0.06)	(0.04)	(0.05)	(0.04)	(0.04)	(0.04)
Auto	0.50***	0.20***	0.30***	0.16***	-0.02	0.18***	-0.01	-0.14***	0.13***
	(0.11)	(0.07)	(0.08)	(0.04)	(0.03)	(0.04)	(0.02)	(0.02)	(0.02)
No-Face	-0.81***	-0.55***	-0.26**	-0.30***	-0.21***	-0.10	-0.03	-0.02	-0.01
	(0.13)	(0.09)	(0.09)	(0.04)	(0.04)	(0.03)	(0.04)	(0.02)	(0.01)
Not Onsite	1.38***	1.47***	-0.10	0.22***	0.31***	-0.09**	0.09*	0.06*	0.03
	(0.18)	(0.17)	(0.15)	(0.03)	(0.03)	(0.02)	(0.05)	(0.03)	(0.02)
No Decision	-0.70***	-0.50***	-0.19***	0.12***	0.13***	-0.01	0.27***	0.28***	-0.01
	(0.13)	(0.09)	(0.06)	(0.03)	(0.03)	(0.02)	(0.05)	(0.04)	(0.04)
2) Detailed Wage Structure Effects									
Union	-1.27***	0.80*	-2.07***	0.75***	0.24*	0.52***	0.35*	0.71***	-0.36*
	(0.35)	(0.41)	(0.39)	(0.16)	(0.14)	(0.17)	(0.18)	(0.15)	(0.22)
Education	4.98***	2.66**	2.32	6.07***	0.58	5.49***	2.28**	0.62	1.66*
	(1.64)	(1.19)	(1.55)	(0.77)	(0.56)	(0.77)	(0.95)	(0.66)	(0.89)
Experience	3.79	1.81	1.98	-1.78	-0.67	-1.11	-4.08***	0.01	-4.09***
	(2.48)	(1.75)	(2.08)	(1.28)	(0.85)	(1.16)	(1.45)	(1.02)	(1.46)
Info	0.94	-2.28**	3.21***	1.99***	1.00**	0.99**	-1.20*	1.48***	-2.68***
	(0.99)	(0.93)	(1.05)	(0.58)	(0.43)	(0.47)	(0.64)	(0.53)	(0.59)
Auto	3.20***	1.44**	1.76**	0.48	-2.18***	2.67***	0.55	-0.28	0.84
	(0.85)	(0.72)	(0.85)	(0.35)	(0.32)	(0.37)	(0.53)	(0.4)	(0.56)
No-Face	-1.10	-1.99**	0.89	1.05**	0.07	0.99**	-1.25*	1.20***	-2.45***
	(1.25)	(0.88)	(0.91)	(0.52)	(0.36)	(0.49)	(0.67)	(0.44)	(0.67)
Not Onsite	4.96***	1.06	3.90***	1.31***	-0.55	1.86***	1.84***	-0.30	2.14***
	(1.08)	(0.9)	(0.86)	(0.48)	(0.4)	(0.42)	(0.54)	(0.4)	(0.51)
No Decision	0.19	1.31	-1.12	-0.33	-0.08	-0.25	-3.27***	-3.62***	0.34
	(1.36)	(1.09)	(1.12)	(0.61)	(0.43)	(0.55)	(0.68)	(0.5)	(0.71)

Note: Log wage differentials × 100. Bootstrapped standard errors are in parentheses (100 replications of the entire

Appendix TableA5. Detailed Decomposition Results - Main Factors: Women

Years	A: 1976/78 to 1988/90			B: 1988/90 to 2000/02			C: 2000/02 to 2011/12		
Inequality Measure:	90-10	50-10	90-50	90-10	50-10	90-50	90-10	50-10	90-50
1) Detailed Composition Effects:									
Union	-0.45***	-0.69***	0.24***	-0.01	-0.13***	0.13***	0.03***	-0.13***	0.16***
	(0.05)	(0.07)	(0.05)	(0.01)	(0.26)	(0.03)	(0.01)	(0.02)	(0.03)
Education	2.80***	2.21***	0.59***	2.93***	2.96***	-0.03	4.83***	3.00***	1.83***
	(0.21)	(0.15)	(0.17)	(0.14)	(0.09)	(0.11)	(0.17)	(0.1)	(0.15)
Experience	1.39***	1.15***	0.24**	0.29***	0.44***	-0.15**	-0.25***	0.17***	-0.42***
	(0.11)	(0.09)	(0.12)	(0.06)	(0.04)	(0.06)	(0.06)	(0.04)	(0.06)
Info	-0.55***	0.97***	-1.52***	0.06	0.58***	-0.52***	-0.04**	-0.11***	0.07***
	(0.18)	(0.16)	(0.26)	(0.06)	(0.06)	(0.06)	(0.02)	(0.02)	(0.02)
Auto	0.22	0.09	0.13	0.73***	0.28***	0.45***	0.35***	0.13***	0.21***
	(0.24)	(0.1)	(0.14)	(0.2)	(0.08)	(0.13)	(0.05)	(0.02)	(0.03)
No-Face	-0.04	-0.08***	0.04	0.03	-0.01	0.04	-0.02	-0.01	-0.01
	(0.03)	(0.03)	(0.03)	(0.02)	(0.01)	(0.02)	(0.04)	(0.01)	(0.02)
Not Onsite	0.35***	0.39***	-0.04	0.38***	-0.08***	0.46***	-0.08***	0.00	-0.08***
	(0.1)	(0.08)	(0.09)	(0.04)	(0.03)	(0.04)	(0.02)	(0.01)	(0.02)
No Decision	-2.24***	-2.85***	0.62*	0.07	0.33***	-0.26***	0.00	0.40***	-0.40***
	(0.33)	(0.42)	(0.33)	(0.06)	(0.05)	(0.04)	(0.05)	(0.04)	(0.05)
2) Detailed Wage Structure Effects									
Union	-1.23***	-0.59***	-0.64	-0.37**	-0.43***	0.07	0.19	0.17	0.02
	(0.32)	(0.2)	(0.39)	(0.18)	(0.1)	(0.19)	(0.21)	(0.12)	(0.22)
Education	5.49***	3.62***	1.87	7.38***	1.68***	5.70***	2.90***	-1.82***	4.72***
	(1.52)	(1.1)	(1.39)	(0.67)	(0.52)	(0.6)	(0.88)	(0.57)	(0.87)
Experience	0.55	-0.64	1.19	-1.52	-0.35	-1.17	-2.55*	0.95	-3.50**
	(2.73)	(1.55)	(2.7)	(1.18)	(0.75)	(1.04)	(1.4)	(0.85)	(1.36)
Info	-1.41	-2.04**	0.63	2.06***	-0.81*	2.88***	7.60***	6.09***	1.50**
	(1.75)	(1.02)	(1.45)	(0.57)	(0.46)	(0.5)	(0.68)	(0.42)	(0.64)
Auto	4.60***	3.90***	0.70	0.31	0.65**	-0.35	1.76***	1.60***	0.16
	(1.13)	(0.84)	(0.91)	(0.46)	(0.3)	(0.4)	(0.59)	(0.34)	(0.54)
No-Face	-1.79*	1.97***	-3.77***	2.19***	0.55**	1.64***	0.32	-0.19	0.51
	(1.06)	(0.65)	(1.07)	(0.43)	(0.21)	(0.42)	(0.62)	(0.28)	(0.61)
Not Onsite	-3.59***	0.42	-4.02***	-0.54*	1.00***	-1.53***	-0.59	1.63***	-2.22***
	(1.17)	(0.7)	(1.08)	(0.36)	(0.22)	(0.37)	(0.54)	(0.3)	(0.49)
No Decision	-0.46	2.70**	-3.16*	-1.31***	-1.43***	0.12	-1.33**	-1.20***	-0.13
	(2.01)	(1.24)	(1.63)	(0.44)	(0.32)	(0.41)	(0.67)	(0.37)	(0.63)

Note: Log wage differentials × 100. Bootstrapped standard errors are in parentheses (100 replications of the entire procedure). Explanatory variables include dummies for union coverage, married, non-white, 6 education classes, 9 experience classes, 4 quartiles of the task indexes. The reported effect for the task indexes is for the upper quartile when the bottom quartile is omitted. Asterisks indicate statistical significance at the 1% (***), 5% (**) or 10%(*) level.

Appendix Figure A1. Occupational Wage Differences With and Without Selection Adjustment (Fixed Effect): PSID 1976-2007

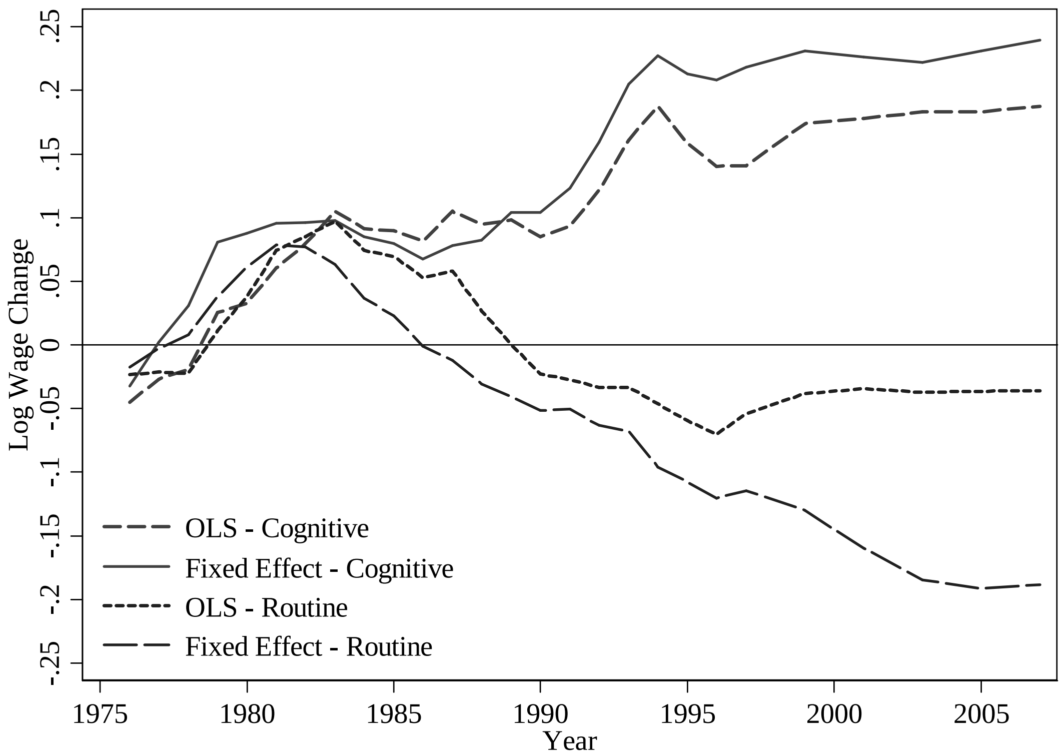

Appendix Figure A2. Cyclicality in Measures of Wage Dispersion: 1973-2012

A. Men

B. Women

Appendix Figure A3. Coefficients of RIF-Regressions – Men
Selected Variables and Years

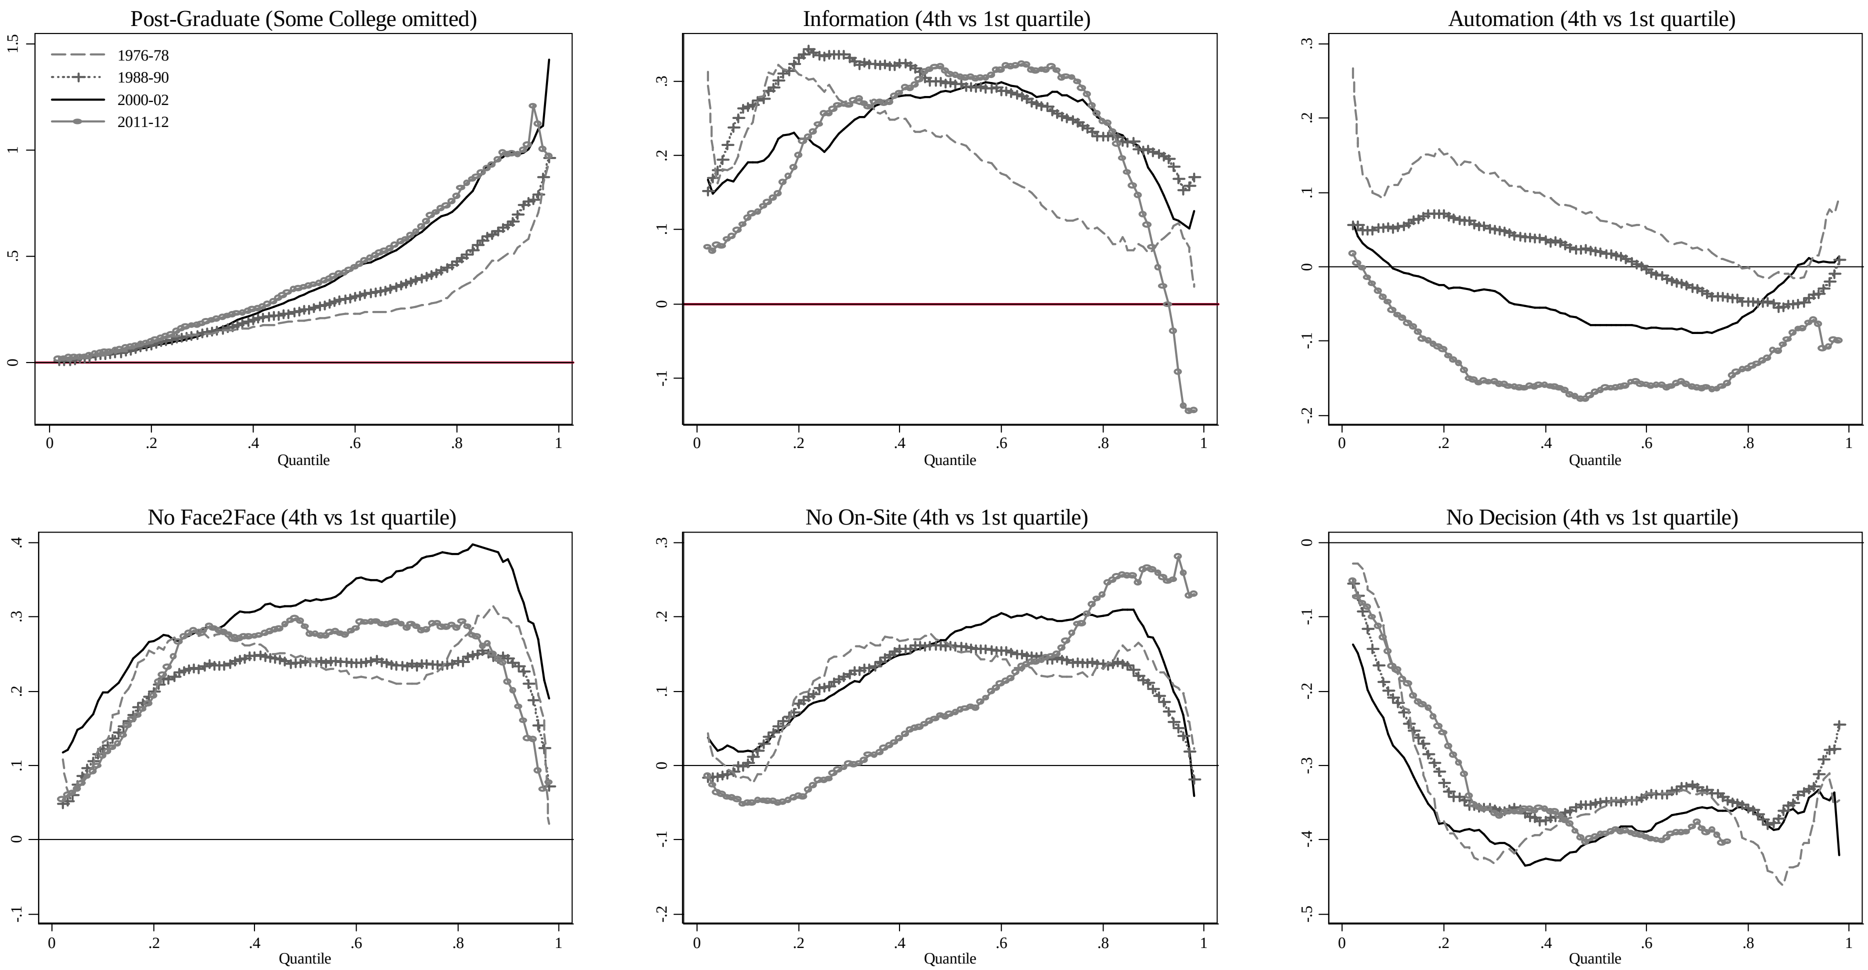

Appendix Figure A4. Coefficients of RIF-Regressions – Women
Selected Variables and Years

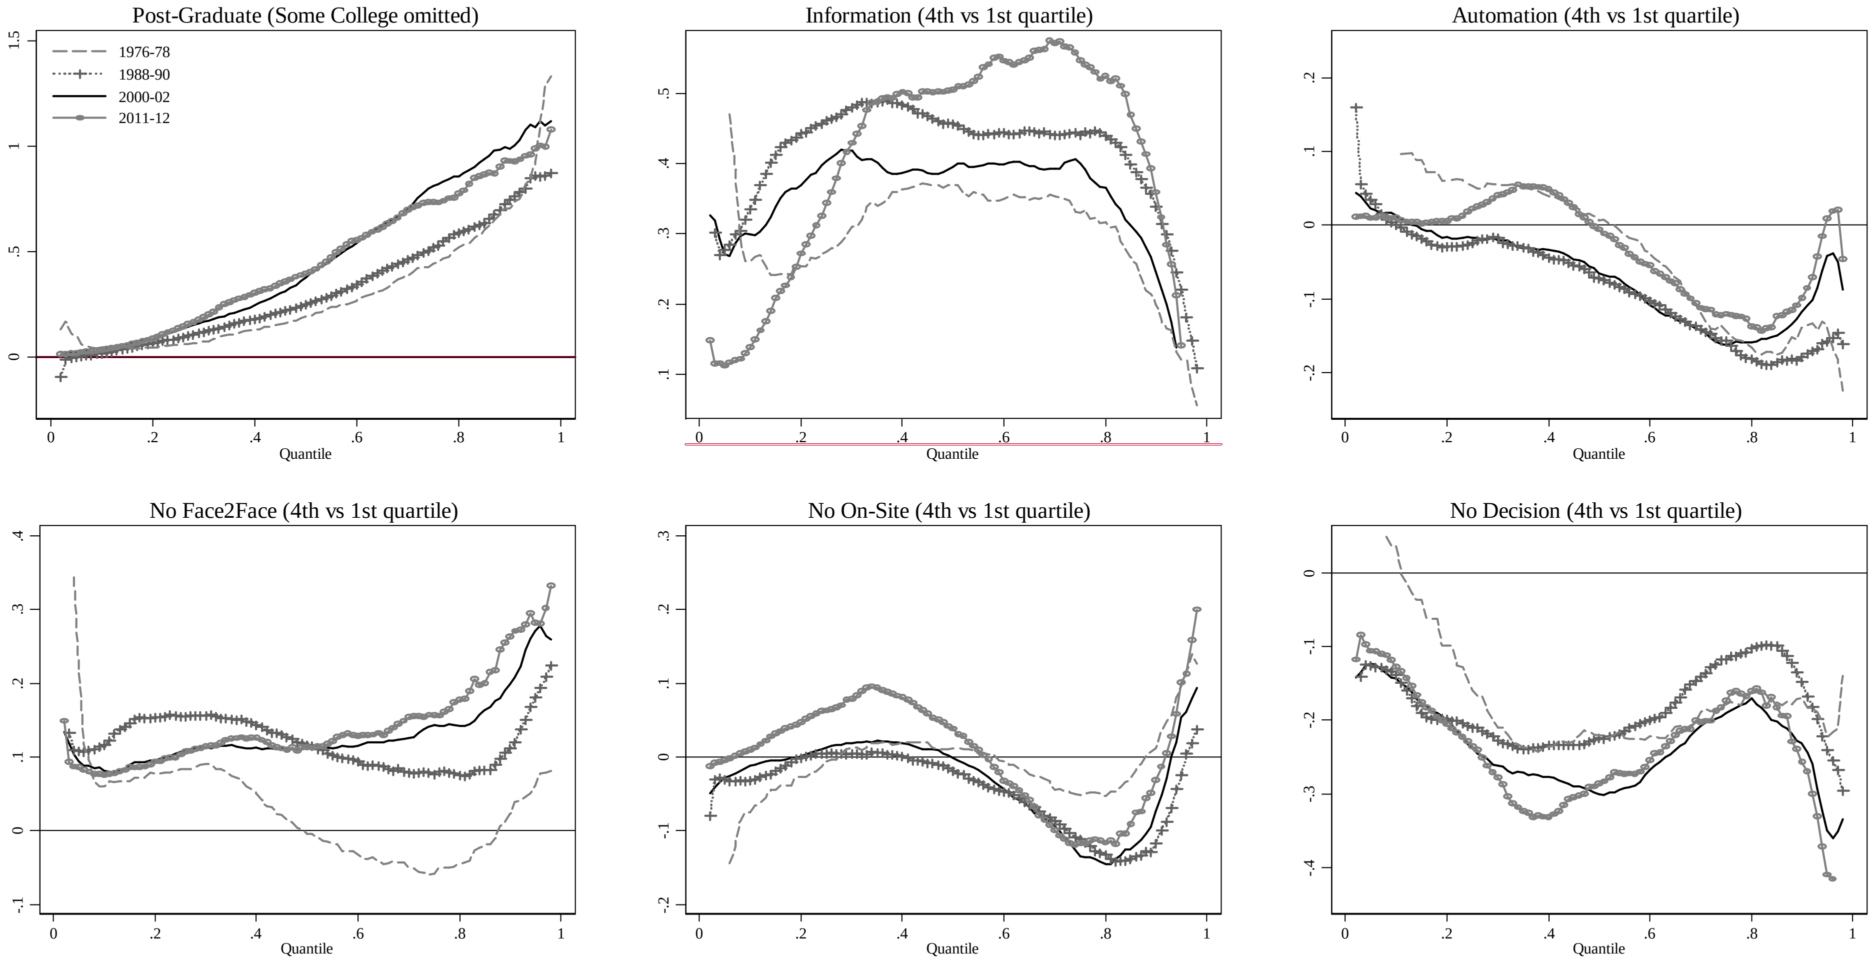

Appendix Figure A5. Aggregate Decomposition Results

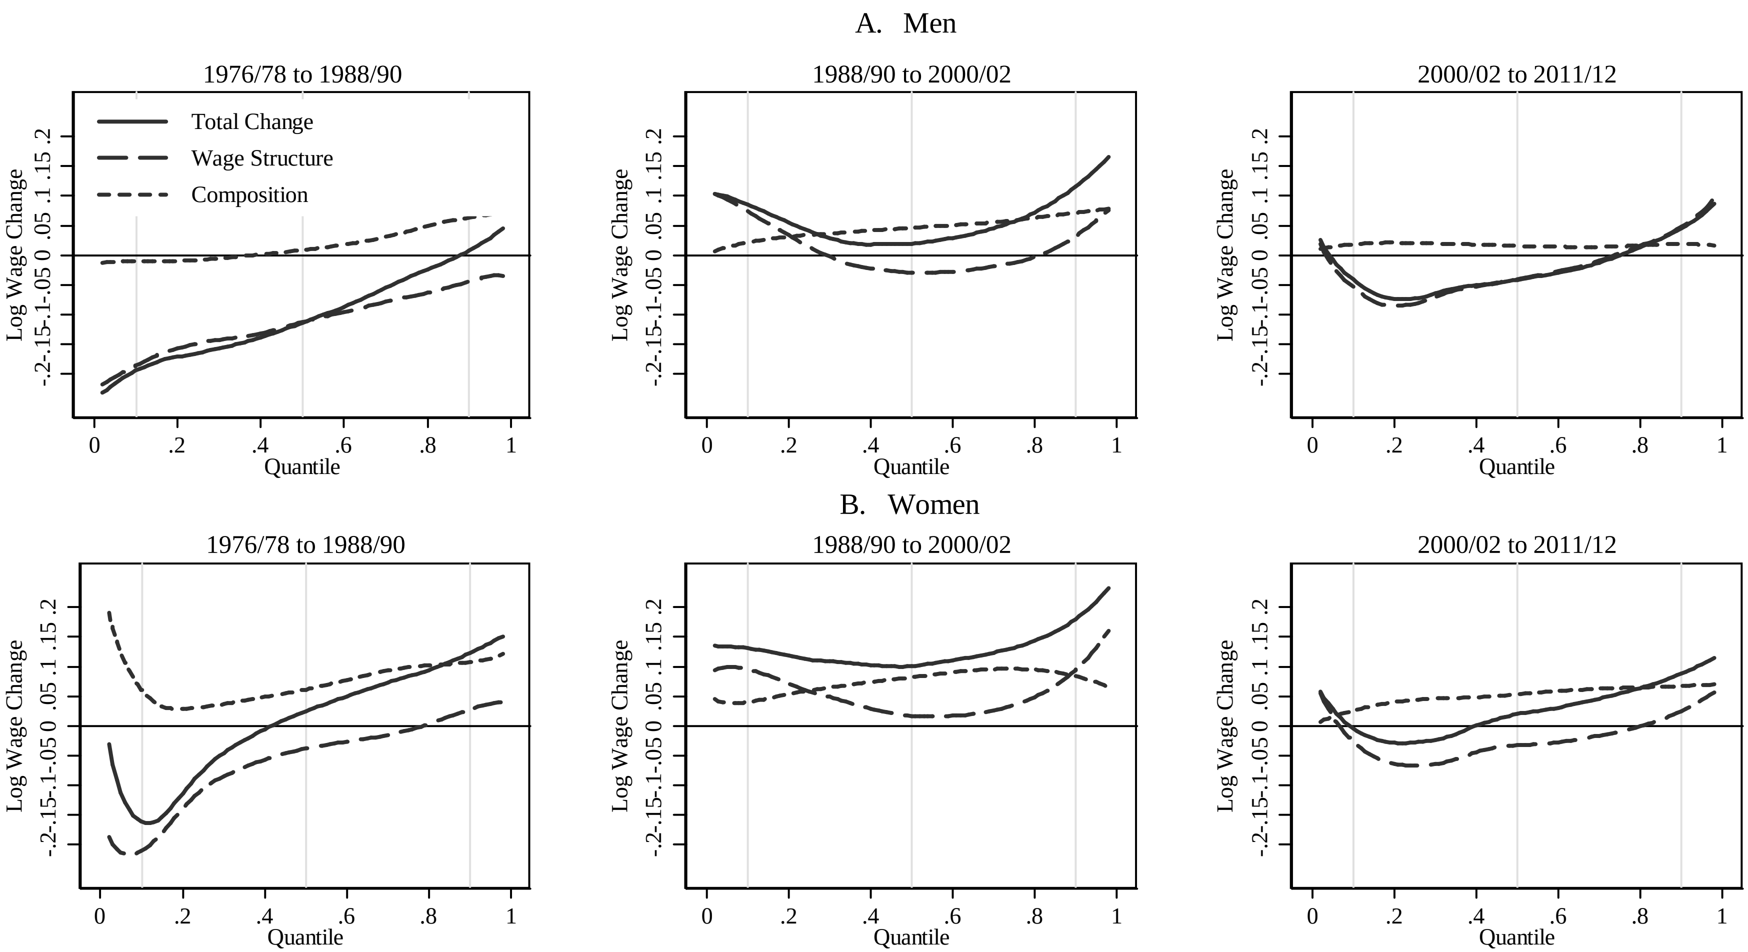

Appendix Figure A6. Detailed Decomposition of Wage Structure Effects
Alternate Grouping of O*Net Measures

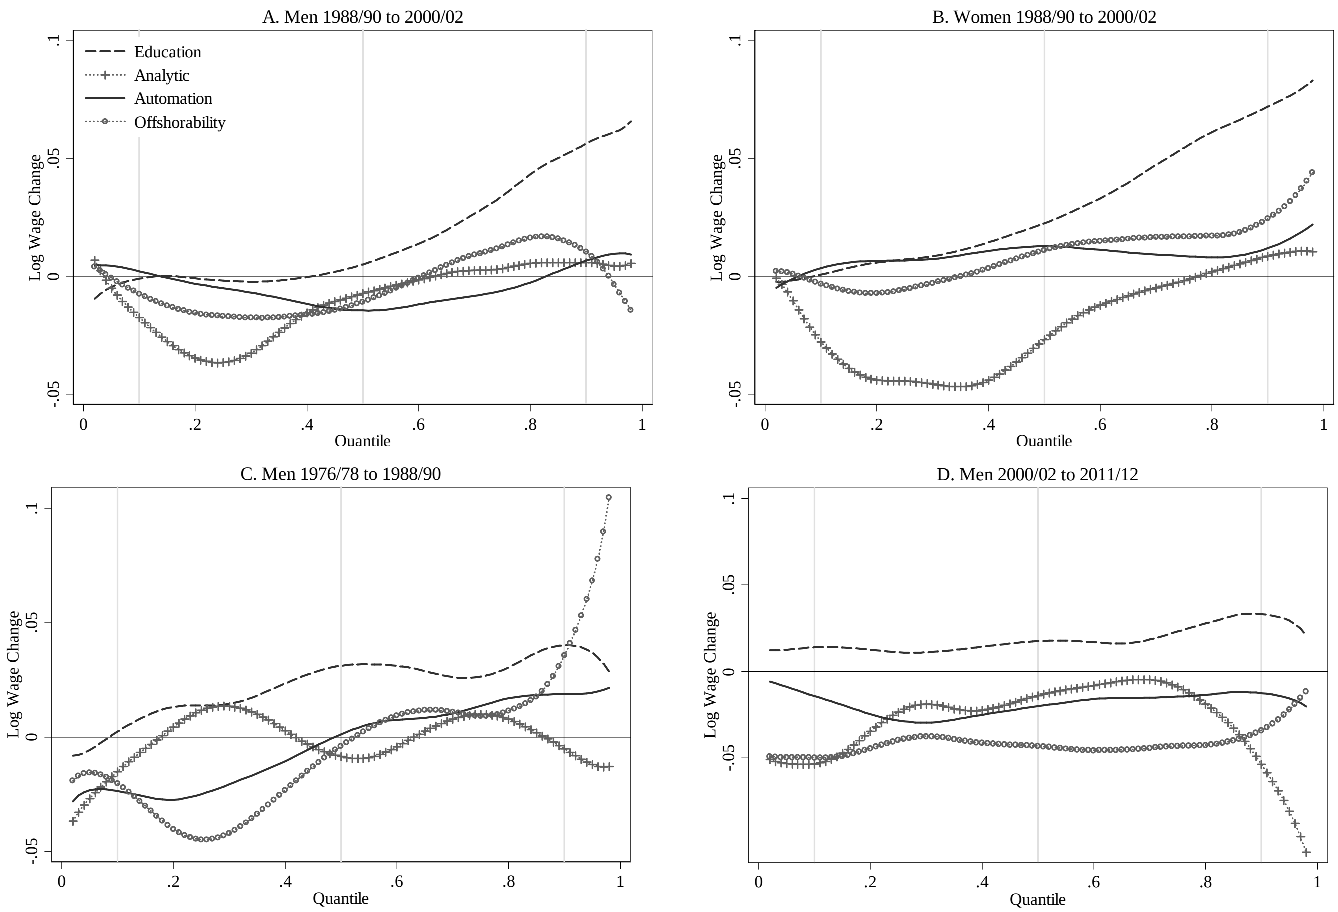

Note: These measures combine the five task content measures (top quartile) in the following way: Analytic is the sum of Information and Decision-making, Offshorability is the sum of No-face-to-face and No On-Site, Automation is simply the measure of automation described in Table A2.

Appendix B - Connection between the Wage Setting Model and RIF-regressions

B-1. General Setup

In Section III, we propose using RIF-regressions to estimate the impact of various factors at different points of the wage distribution. In this Appendix we establish the connection between this decomposition method and the wage setting model of Section I. We begin with a general model and then look at a special case to help with interpretation.

Consider a very general model where the wage of worker i can be written as

$$w_i = m\left(O_i, S_i; \theta\right), \tag{B-1}$$

where O is a categorical variable indicating occupation type (for instance, working with the three digit occupational classification we would have $O = 55$ for electrical engineers, $O = 567$ for carpenters, etc.); θ is a parameter vector; the amount of tasks produced by worker i depends on a single skill index $S_i = \sum_{k=1}^{K} \alpha_k S_{ik}$. We omit time subscripts to simplify the notation and assume m is increasing in S. Since some of the K components of S are not observed, we simply consider S as being the unobservable component. Equation (B-1) is general enough to accommodate any particular wage setting model where skills are used to produce different tasks in different occupations, and workers are paid on the basis of how much tasks they produce.

Consider for example what happens to the wage distribution when routine tasks become less valuable in the labor market because of an exogenous technological change. This corresponds in this very general setup to a change in one of the entries of θ. We first present calculations for general changes in θ and then present in the next subsection a special case of equation (4) in which one of the element of θ corresponds to the return to routine tasks. Readers mainly interested in this application may want to skip to the next subsection.

Let q_τ represent the τ^{th} quantile of the wage distribution, that is, $q_\tau = F_w^{-1}(\tau)$, where $F_w(.)$ is the CDF of wages. Since

$$\begin{aligned}\tau &= F_w(q_\tau) = \Pr\left[m\left(O, S; \theta\right) \leq q_\tau\right] = \Pr\left[S \leq m^{-1}\left(O, q_\tau; \theta\right)\right] \\ &= E\left[\Pr\left[S \leq m^{-1}\left(O, q_\tau; \theta\right)|O\right]\right] = E\left[F_{S|O}\left(m^{-1}\left(O, q_\tau; \theta\right)|O\right)\right]\end{aligned}$$

we can apply the Implicit Function Theorem to write the effect of changes in θ on q_τ as

$$
\begin{aligned}
\frac{\partial q_\tau}{\partial \theta} &= -\frac{\partial E\left[F_{S|O}\left(m^{-1}\left(O, q_\tau; \theta\right)|O\right)\right]/\partial\theta}{f_w(q_\tau)} \\
&= \frac{\sum_{j=1}^{J} f_{S|O}\left(m^{-1}\left(O, q_\tau; \theta\right)|O = O_j\right) \frac{m_\theta\left(O_j, m^{-1}(O_j, q_\tau; \theta); \theta\right)}{m_S\left(O_j, m^{-1}(O_j, q_\tau; \theta); \theta\right)} P_j}{f_w\left(q_\tau\right)} \\
\text{(B-2)} \qquad &= \frac{\sum_{j=1}^{J} f_{S|O}\left(m^{-1}\left(O, q_\tau; \theta\right)|O = O_j\right) \frac{m_\theta\left(O_j, m^{-1}(O_j, q_\tau; \theta); \theta\right)}{m_S\left(O_j, m^{-1}(O_j, q_\tau; \theta); \theta\right)} P_j}{\sum_{j=1}^{J} f_{S|O}\left(m^{-1}\left(O, q_\tau; \theta\right)|O = O_j\right)\left(m_S\left(O_j, m^{-1}\left(O_j, q_\tau; \theta\right); \theta\right)\right)^{-1} P_j}
\end{aligned}
$$

where $f_{S|O}$ is the density of skill S in occupation O, f_w is the marginal wage density, $P_j = \Pr\left[O = O_j\right]$ is the probability of being in occupation O_j, and m_θ and m_S are derivatives of m with respect to θ and skill S, respectively.

Now consider the effect of changes in the probability of being in an occupation O_l on q_τ. That is, consider $\partial q_\tau/\partial P_l$. Firpo, Fortin and Lemieux (2009) call $\partial q_\tau/\partial P_l$ an "unconditional quantile partial effect" (UQPE), and show that it can be estimated using a RIF-Regression for the τ^{th} quantile. The estimate of $\partial q_\tau/\partial P_l$ is the coefficient on a dummy variable for occupation l in a RIF-Regression like the one introduced in Section III.

Using the Implicit Function Theorem, we can write the UQPE $\partial q_\tau/\partial P_l$ as:

$$
\gamma_{\tau l} = \frac{\partial q_\tau}{\partial P_l} = \frac{\Pr\left[S \geq m^{-1}\left(O, q_\tau; \theta\right)|O = O_l\right] - \Pr\left[S \geq m^{-1}\left(O, q_\tau; \theta\right)|O = o\right]}{f_w(q_\tau)},
$$

where occupation $O = o$ is the baseline occupation.

A marginal change in θ has the following effect on $\gamma_{\tau l}$:

$$
\begin{aligned}
\frac{\partial \gamma_{\tau l}}{\partial \theta} &= \Bigg(f_{S|O}\left(m^{-1}\left(O, q_\tau; \theta\right)|O = O_l\right)\left(\frac{m_\theta\left(O_l, m^{-1}\left(O_l, q_\tau; \theta\right); \theta\right)}{m_S\left(O_l, m^{-1}\left(O_j, q_\tau; \theta\right); \theta\right)}\right) \\
&\quad - f_{S|O}\left(m^{-1}\left(O, q_\tau; \theta\right)|O = o\right)\left(\frac{m_\theta\left(o, m^{-1}\left(o, q_\tau; \theta\right); \theta\right)}{m_S\left(o, m^{-1}\left(o, q_\tau; \theta\right); \theta\right)}\right)\Bigg) / f_w(q_\tau).
\end{aligned}
$$

In order to simplify our calculations, let's normalize the derivative in the baseline occupation to zero ($m_\theta\left(o, m^{-1}\left(o, q_\tau; \theta\right); \theta\right) = 0$).[46] We can then see that $\partial\gamma_{\tau l}/\partial\theta$ and $\partial q_\tau/\partial\theta$ are closely linked:

[46]In a standard linear regression model, the parameters θ are the coefficients on the occupational dummies relative to the base occupation. Under this conventional normalization (include dummies for all occupations except the base), a change in θ has no effect on wages in the base occupation. In other words, the derivative $m_\theta(.)$ is equal to zero for the base occupation. The normalization we use here is simply a generalization of this conventional normalization in the context of our general wage setting model $m\left(O_i, S_i; \theta\right)$.

$$\frac{\partial q_\tau}{\partial \theta} = \sum_{j=1}^{J} \frac{\partial \gamma_{\tau j}}{\partial \theta} P_j = E\left[\frac{\partial \gamma_\tau}{\partial \theta}\right]. \tag{B-3}$$

This important result indicates that the effect of changes in θ on q_τ will be reflected in changes in the RIF-Regression coefficients γ_τ. To make this result more concrete, consider changes over time in the wage distribution, since this is the main object of interest in this paper. Let the structural parameter θ assume value θ_0 at time 0 and θ_1 at time 1. Similarly, consider the RIF-Regression coefficients of the τ^{th} quantile on a dummy variable for occupation l for time periods 1 and 0. Using equation (B-3), it follows that:

$$\gamma_{\tau l}\big|_{\theta=\theta_1} - \gamma_{\tau l}\big|_{\theta=\theta_0} \approx \left.\frac{\partial \gamma_l}{\partial \theta}\right|_{\theta=\theta_0} (\theta_1 - \theta_0),$$

and therefore

$$\begin{aligned} q_\tau(\theta_1) - q_\tau(\theta_0) &\approx \left.\frac{\partial q_\tau}{\partial \theta}\right|_{\theta=\theta_0} (\theta_1 - \theta_0) \\ &= \sum_{j=1}^{J} \left(\left.\frac{\partial \gamma_{\tau j}}{\partial \theta}\right|_{\theta=\theta_0} P_j \right) (\theta_1 - \theta_0) \\ &\approx \sum_{j=1}^{J} \left(\gamma_{\tau j}\big|_{\theta=\theta_1} - \gamma_{\tau j}\big|_{\theta=\theta_0} \right) P_j. \qquad \text{(B-4)} \end{aligned}$$

This equation illustrates the close connection between the effect of changes in the structural parameters θ on q_τ over time and the changes in the RIF-Regression coefficients. The last component of equation (B-4) corresponds to the wage structure component of our decomposition presented in equation (9) of Section III where changes over time in the RIF-Regression coefficients are multiplied by the average values of the covariates (average fraction of workers in each occupation in the case considered here). Importantly, this suggests that we can estimate the effect of changes in the structural parameters θ by simply performing a Oaxaca-type decomposition using RIF-Regressions.

The result obtained in equation (B-4) is very general. It indicates that the impact of changes over time in the whole vector of structural parameters θ can be estimated using the wage structure component of a generalized Oaxaca-type decomposition based using RIF-Regressions. In this general setting it is not possible, however, to know how changes in a particular element of θ (e.g.

the return to routine tasks) maps into changes in a particular RIF-Regression coefficient (like the one associated to a dummy variable indicating whether occupation j is a routine occupation). We next look at a special case based on our main wage setting equation (4) where it is possible to establish this connection.

B-2. An example

Consider a simple case with three occupations: a base –or unskilled– occupation o, a "routine" occupation r, and an "offshorable" occupation a. Recall that in our main decomposition we use a set of dummy variables to capture the different quartiles of the distribution of the task content measures. Here we further simplify the approach by using single dummy variables to indicate whether the occupation is routine or offshorable. In other words, the routine occupation has a routine task score, T_r, of one, while the other occupations have routine task scores of zero. Likewise, the offshorable occupation has an offshorability task score, T_a, of one, while the other occupations have offshorability task scores of zero. Using the wage setting model of equation (4), we get:

$$w_i = c_0 + \pi_0 S_i + T_{ir}(c_r + \pi_r S_i) + T_{ia}(c_a + \pi_a S_i). \tag{B-5}$$

In this model, in the baseline occupation we have $T_{ir} = T_{ia} = 0$ and $w_i = c_0 + \pi_0 S_i$; in the routine-based occupation, $T_{ir} = 1 - T_{ia} = 1$ and $w_i = c_0 + c_r + (\pi_0 + \pi_r) S_i$; and in the offshorable occupation, $1 - T_{ir} = T_{ia} = 1$ and $w_i = c_0 + c_a + (\pi_0 + \pi_a) S_i$. The parameter vector θ in this case is $[c_0, c_r, c_a, \pi_0, \pi_r, \pi_a]^\top$.

Now consider $\partial q_\tau / \partial \pi_r$, the effect of a change in the return to routine tasks π_r on the τ^{th} quantile (q_τ) of the wage distribution. $\partial q_\tau / \partial \pi_r$ can be written down explicitly as a function of the structural parameters of the model using equation (B-2). In the special case considered here, the skill density for each occupation O at $m^{-1}(O, q_\tau; \theta)$ is

$$f_{S|O}\left(m^{-1}(O, q_\tau; \theta) | O\right) = \begin{cases} f_0\left((q_\tau - c_0)/\pi_0\right), \text{ for the baseline occupation} \\ f_r\left((q_\tau - c_0 - c_r)/(\pi_0 + \pi_r)\right), \text{ for routine occupation} \\ f_a\left((q_\tau - c_0 - c_a)/(\pi_0 + \pi_a)\right), \text{ for the offshorable occupation} \end{cases}$$

and $f_w(q_\tau)$, the wage density at q_τ, is

$$\begin{aligned} f_w(q_\tau) &= f_0\left(\left(q_\tau - c_0\right)/\pi_0\right)\left(\left(1 - P_r - P_a\right)/\pi_0\right) \\ &\quad + f_r\left(\left(q_\tau - c_0 - c_r\right)/\left(\pi_0 + \pi_r\right)\right)\left(P_r/\left(\pi_0 + \pi_r\right)\right) \\ &\quad + f_a\left(\left(q_\tau - c_0 - c_a\right)/\left(\pi_0 + \pi_a\right)\right)\left(P_a/\left(\pi_0 + \pi_a\right)\right). \end{aligned}$$

Likewise, the derivative of m with respect to skill evaluated at skill level $m^{-1}(O, q_\tau; \theta)$ is:

$$m_S\left(O, m^{-1}\left(O, q_\tau; \theta\right); \theta\right) = \begin{cases} \pi_0, \text{ in the baseline occupation} \\ \pi_0 + \pi_r, \text{ in the routine occupation} \\ \pi_0 + \pi_a, \text{ in the offshorable occupation} \end{cases}$$

while $m_{\pi_r}\left(O, m^{-1}\left(O, q_\tau; \theta\right); \theta\right) = \left(q_\tau - c_0 - c_r\right)/\left(\pi_0 + \pi_r\right)$ if O is the routine occupation, and $m_{\pi_r}\left(O, m^{-1}\left(O, q_\tau; \theta\right); \theta\right) = 0$ otherwise. We also have $P_r = \Pr\left[O = \text{routine occupation}\right]$ and $P_a = \Pr\left[O = \text{offshorable occupation}\right]$ denote the fraction of the workforce in the routine and offshorable occupations, respectively.

Substituting all these expressions in equation (B-2) we get

(B-6) $$\frac{\partial q_\tau}{\partial \pi_r} = \frac{f_r\left(\left(q_\tau - c_0 - c_r\right)/\left(\pi_0 + \pi_r\right)\right)\left(q_\tau - c_0 - c_r\right)/\left(\pi_0 + \pi_r\right)^2 P_r}{f_w\left(q_\tau\right)}.$$

Equation (B-6) is easier to interpret by noting that the term $\left(q_\tau - c_0 - c_r\right)/\left(\pi_0 + \pi_r\right)$ represents the skill level that a worker in the routine occupation needs to earn the wage q_τ. Call this skill level $S(r, q_\tau)$, where $S(r, q_\tau) = m^{-1}\left(O = r, q_\tau; \theta\right) = \left(q_\tau - c_0 - c_r\right)/\left(\pi_0 + \pi_r\right)$.

Furthermore, we can replace the skill density $f_r(.)$ by a wage density using the standard formula

$$f_{w|O}\left(q_\tau|O\right) = \frac{f_{S|O}\left(m^{-1}\left(O, q_\tau; \theta\right)|O\right)}{m_S\left(O, m^{-1}\left(O, q_\tau; \theta\right); \theta\right)}.$$

Let $f_{wr}\left(q_\tau\right) = f_{w|O}\left(q_\tau|O = r\right) = f_r\left(\left(q_\tau - c_0 - c_r\right)/\left(\pi_0 + \pi_r\right)\right)/\left(\pi_0 + \pi_r\right)$ be the wage density in the routine occupation. It follows that

(B-7) $$\frac{\partial q_\tau}{\partial \pi_r} = P_r \frac{f_{wr}\left(q_\tau\right)}{f_w(q_\tau)} S(r, q_\tau).$$

This equation shows that a change in π_r can have quite different impacts at different points of the wage distribution. In particular, a change in π_r will have a larger impact in the part of the distribution where routine workers are concentrated, i.e. where the relative wage density $f_{wr}(q_\tau)/f_w(q_\tau)$ is the highest. This is consistent with Autor, Katz, and Kearney (2006) who point out that if "routine workers" are concentrated in the middle of the distribution, we should expect technological change –which depresses the return to these routine tasks– to have its largest impact in that part of the distribution. This illustrates the general point made in Section I that our wage setting model has complex implications for the wage distribution because task prices affect both the between- and within-group component of wage inequality. Equation (B-7) also illustrates the importance of using a flexible approach such as RIF-Regressions to capture different effects at different point of the wage distribution.

In the general model developed earlier, we were not able to show a direct correspondence between a specific structural parameter and a specific RIF-Regression coefficient. We are now able to do so in the special case considered here. Using equation (B-3), we have

$$\begin{aligned}
\frac{\partial q_\tau}{\partial \pi_r} &= \sum_{j=1}^{J} \frac{\partial \gamma_{\tau j}}{\partial \pi_r} P_j \\
&= \frac{1}{f_w(q_\tau)} \sum_{j=1}^{J} P_j \left(f_{w|O}(q_\tau|O_j)\, m_{\pi_r}\left(O_j, m^{-1}(O_l, q_\tau;\theta);\theta\right)\right. \\
&\quad \left. - f_{w|O}(q_\tau|o)\, m_{\pi_r}\left(o, m^{-1}(o, q_\tau;\theta);\theta\right)\right).
\end{aligned}$$

Since $m_{\pi_r}\left(O, m^{-1}(O, q_\tau;\theta);\theta\right) = S(r, q_\tau)$ if O is the routine occupation and zero otherwise, it follows that the derivative of the RIF-Regression coefficient is also zero for all occupations but the routine occupation. As a result:

$$\frac{\partial \gamma_{\tau j}}{\partial \pi_r} = \begin{cases} \frac{f_{wr}(q_\tau)}{f_w(q_\tau)} S(r, q_\tau) \text{ if } j = r \\ 0 \text{ otherwise} \end{cases},$$

and therefore

$$\frac{\partial q_\tau}{\partial \pi_r} = P_r \frac{f_{wr}(q_\tau)}{f_w(q_\tau)} S(r, q_\tau) = P_r \frac{\partial \gamma_{\tau r}}{\partial \pi_r}.$$

We can finally use equation (B-4) to show that a change over time in the return to routine tasks (from π_{r0} to π_{r1}) can be estimated using the wage structure component corresponding to the routine occupation:

$$q_\tau\left(\pi_{r1} : \widetilde{\theta}_0\right) - q_\tau\left(\pi_{r0} : \widetilde{\theta}_0\right) \approx \left(\gamma_{\tau r1} - \gamma_{\tau r0}\right) P_r, \tag{B-8}$$

where $\widetilde{\theta}$ is a vector of all the structural parameters except π_r.

Appendix C - Empirical Test of the Occupational Wage Setting Model

C-1. Simple implications for means and standard deviations

In this Appendix we discuss in more detail the empirical evidence in support of the validity of the wage setting model introduced in Section I. To fix ideas, consider a simplified version of equation (2) where we ignore the covariates Z_{it}:

$$w_{ijt} = \delta_t + c_j + p_{jt} Y_{ij} \equiv \delta_t + c_j + p_{jt} \sum_{k=1}^{K} \alpha_{jk} S_{ik}. \tag{C-1}$$

As we discuss in Section I, in this model changes in task prices p_{jt} have an impact on both the level and dispersion of wages across occupations. For instance, the average wage in occupation j at time t is

$$\overline{w}_{jt} = \delta_t + c_j + p_{jt} \overline{Y}_{jt}. \tag{C-2}$$

The standard deviations of wages is

$$\sigma_{jt} = p_{jt} \sigma_{Y,jt}, \tag{C-3}$$

where $\sigma_{Y,jt}$ is the standard deviation in tasks Y_{ij}, which in turns depends on the within-occupation distribution of skills S_{ik}. Since changes in both $\overline{w}_{jt}$ and σ_{jt} are positively related to changes in task prices p_{jt}, we expect these two changes to be correlated across occupations.

To see this more formally, assume that the within-occupation distribution of skills, S, and thus the distribution of task output, Y, remains constant over time (we discuss the assumption in more detail below). It follows that $\overline{Y}_{jt} = \overline{Y}_j$ and $\sigma_{Y,jt} = \sigma_{Y,j}$ for all t. Using a first order approximation of equations (C-2) and (C-3) and differencing yields:

$$\Delta \overline{w}_j \approx \Delta \delta + \overline{Y} \cdot \Delta p_j, \tag{C-4}$$

and

$$\Delta \sigma_j \approx \overline{\sigma_Y} \cdot \Delta p_j, \tag{C-5}$$

where $\overline{Y}$ ($\overline{\sigma_Y}$) is the average of $\overline{Y}_j$ ($\sigma_{Y,j}$) over all occupations j. Since the variation in Δp_j is the only source of variation in $\Delta\overline{w}_j$ and $\Delta\sigma_j$, the correlation between these two variables should be equal to one in this simplified model. In practice, we expect the correlation to be fairly large and positive, but not quite equal to one because of sampling error (in the estimates values of $\Delta\overline{w}_j$ and $\Delta\sigma_j$), approximation errors, etc.

A second implication of the model is that since task prices p_{jt} depend on the task content measures T_{jh} (see equation 3), these tasks content measures should help predict changes in task prices Δp_j, and thus $\Delta\overline{w}_j$ and $\Delta\sigma_j$. Differencing equation (3) over time we get:

$$\Delta p_j = \Delta\pi_0 + \sum_{h=1}^{5} \Delta\pi_h T_{jh} + \Delta\mu_j, \tag{C-6}$$

and, thus:

$$\Delta\overline{w}_j = \varphi_{w,0} + \sum_{h=1}^{5} \varphi_{w,h} T_{jh} + \xi_{w,h}, \tag{C-7}$$

and

$$\Delta\sigma_j = \varphi_{\sigma,0} + \sum_{h=1}^{5} \varphi_{\sigma,h} T_{jh} + \xi_{\sigma,h}, \tag{C-8}$$

where $\varphi_{w,0} = \Delta\delta + \overline{Y}\cdot\Delta\pi_0$; $\varphi_{w,h} = \overline{Y}\cdot\Delta\pi_h$; $\xi_{w,h} = \overline{Y}\cdot\Delta\mu_j$; $\varphi_{\sigma,0} = \overline{\sigma_Y}\cdot\Delta\pi_0$; $\varphi_{\sigma,h} = \overline{\sigma_Y}\cdot\Delta\pi_h$; $\xi_{\sigma,h} = \overline{\sigma_Y}\cdot\Delta\mu_j$. One important implication of the model highlighted here is that the coefficients $\varphi_{w,h}$ and $\varphi_{\sigma,h}$ should be proportional in equations (C-7) and (C-8) since they both depend on the same underlying coefficients $\Delta\pi_h$.

C-2. Empirical evidence

We provide evidence that these two implications are supported in the data in the case of men in the 1990s. This group (and time period) is of particular interest since one of the main goal of this paper is to understand the sources of labor market polarization that was particularly important for that group/time period. Note that, despite our large samples based on three years of pooled CPS data, we are left with a small number of observations in many occupations when we work at the three-digit occupation level. In the analysis presented in this Appendix, we thus focus on occupations classified at the two-digit level (40 occupations) to have a large enough number of observations in each occupation.[47] All the estimates reported in this Appendix (correlations and regression models) are weighted using the proportion of workers in the occupation. More details on the CPS and O*NET data (we use the O*NET to construct the task content measures) are provided in Section II.

As we mention in the text, the raw correlation between the changes in average wages and standard deviations is large and positive (0.44), as expected. It increases to 0.57 when we exclude agricultural occupations (less than three percent of the workforce).

We then run regression models for equations (C-7) and (C-8) using our five O*NET task content measures as explanatory variables. Note that in these regressions we use the continuous task content measures instead of the set of quartile dummies discussed in the paper. Quartile dummies make the results of the full decomposition easier to interpret, but would only leave us with few degrees of freedom here since we are only using 40 observations (one change over time for each 2-digit occupation).

The regression results are reported in columns 1-4 of Appendix Table C1. Columns 1 and 2 show the estimated models for $\Delta\overline{w}_j$ and $\Delta\sigma_j$, respectively, when all five task measure variables are included in the regression. The adjusted R-square of the regressions is equal to 0.52 for both models, indicating that our task content measures capture most of the variation in changes in the level ($\Delta\overline{w}_j$) and dispersion ($\Delta\sigma_j$) of wages over occupations. Since several of the coefficients are imprecisely estimated, we also report in columns 3 and 4 estimates from separate regressions for each task content measure. The task content measures are significant in most cases, and the sign of the coefficient estimates are the same in the models for changes in average wages and standard deviations. This strongly support the prediction of our wage setting models that the estimated effect of the task content measures should be proportional in the models for average wages and standard deviations.

Note also that, in most of the cases, the sign of the coefficients conforms to expectations. As some tasks involving the processing of information may be enhanced by ICT technologies, we would expect a positive relationship between our "information content" task measure and changes in task prices. On the other hand, to the extent that technological change allows firms to replace

[47]Though there is a total of 45 occupations at the two-digit level, we combine five occupations with few observations to similar but larger occupations. Specifically, occupation 43 (farm operators and managers) and 45 (forestry and fishing occupations) are combined with occupation 44 (farm workers and related occupations). Another small occupation (20, sales related occupations) is combined with a larger one (19, sales workers, retail and personal services). Finally two occupations in which very few men work (23, secretaries, stenographers, and typists, and 27, private household service occupations) are combined with two other larger occupations (26, other administrative support, including clerical, and 32, personal services, respectively).

workers performing these types of tasks with computer driven technologies, we would expect a negative effect for the "automation/routine" measure. Although occupations in the middle of the wage distribution may be most vulnerable to technological change, some also involve relatively more "on-site" work (e.g. repairmen) and may, therefore, be less vulnerable to offshoring. We also expect workers in occupations with a high level of "face-to-face" contact, as well as those with a high level of "decision-making", to do relatively well in the presence of offshoring. Since these last three variable capture non-offshorability, they are entered as their reverse in the regression (see the discussion in Section II) and we should expect their effect to be negative.

In columns 3 and 4, all the estimated coefficients are of the expected sign except for the "no onsite" task. This may indicate that the O*NET is not well suited for distinguishing whether a worker has to work on "any site" (i.e. an assembly line worker whose job could be offshored), vs. working on a site in the United States (i.e. a construction worker).

One potential issue with these estimates is that we are only using the raw changes in $\overline{w}_{jt}$ and σ_{jt} that are unadjusted for differences in education and other characteristics. Part of the changes in $\overline{w}_{jt}$ and σ_{jt} may, thus, be due to composition effects or changes in the return to underlying characteristics (like education) that are differently distributed across occupations. To control for these confounding factors, we reweight the data using simple logits to assign the same distribution of characteristics to each of the 40 occupations in the two time periods.[48]

This procedure allows us to relax the assumption that the distribution of skills S is constant over time. Strictly speaking, we can only adjust for observable skills like education and experience. To deal with unobservables, we could then invoke an ignorability assumption to ensure that, conditional on observable skills, the distribution of unobservable skills is constant over time. A more conservative approach is to view the specifications where we control for observable skills as a robustness check.

The results reported in columns 5-8 indeed suggest that the main findings discussed above are robust to controlling for observables. Generally speaking, the estimated coefficients have similar magnitudes and almost never change sign relative to the models reported in column 1-4. Overall, the results presented in this appendix strongly support the predictions of our wage setting model.

[48]We use the set of dummies for education, experience, marital status and race described in Section IV in the logits. The estimates are used to construct reweighting factors that are used to make the distribution of characteristics in each occupation-year the same as in the overall sample for all occupations (and time periods).

Appendix Table C1. Estimated Effect of Task Requirements on Average Wages and Standard Deviations
Men, 1988-90 to 2000-02, 2-digit Occupations

Tasks entered:	Together		Separately		Together		Separately	
Dep. variable:	Average (1)	Std dev (2)	Average (3)	Std dev (4)	Average (5)	Std dev (6)	Average (7)	Std dev (8)
	Raw changes				Reweighted changes			
Information	0.0106	0.0041	0.0198***	0.0163***	0.0081***	0.0058	0.0081	0.0179***
content	(0.0108)	(0.0059)	(0.0079)	(0.0039)	(0.0096)	(0.0055)	(0.0061)	(0.0037)
Automation	-0.0306***	-0.0096	-0.0467***	-0.0226***	-0.0137	-0.0146**	-0.0228***	-0.0245***
/routine	(0.0112)	(0.0062)	(0.0089)	(0.0053)	(0.01)	(0.0057)	(0.0078)	(0.0051)
No on-site work	0.0018	0.0099***	0.0190***	0.0132***	0.0021	0.0077**	0.0118***	0.0118***
	(0.0059)	(0.0033)	(0.0048)	(0.0023)	(0.0052)	(0.003)	(0.0037)	(0.0025)
No face-to-face	-0.0418***	0.0082	-0.0560***	-0.0207	-0.0348**	0.0165**	-0.0320**	-0.0180**
	(0.0148)	(0.0081)	(0.0109)	(0.0069)	(0.0131)	(0.0075)	(0.0088)	(0.007)
No decision	0.0229***	-0.0077	-0.0275***	-0.0189***	0.00227*	-0.0125**	-0.0079	0.0221***
making	(0.0163)	(0.009)	(0.0105)	(0.0056)	(0.0144)	(0.0083)	(0.0084)	(0.0052)
Adj. R-square	0.517	0.523	---	---	0.306	0.582	---	---

Notes: All models are estimated by running regressions of the occupation-specific changes in average wages and standard deviations on the task content measures. The models reported in all columns are weighted using the fraction of observations in each occupation in the base period (1988-90). In columns 5-8 the data are reweighted so that the distribution of characteristics in each occupation and time period is the same as in the overall sample (for both periods pooled). See the text for more detail.