

Offshoring and the Polarization of the U.S. Labor Market*

Lindsay Oldenski
Georgetown University

December 2012

Abstract

I use firm-level data on offshoring paired with occupation-level data on employment and wages to estimate the impact that offshoring has had on U.S. workers from 2002 to 2008. I find that offshoring by U.S. firms has contributed to relative gains for the most highly skilled workers and relative losses for middle skilled workers. An increase in offshoring in an industry is associated with an increase in the wage gap between workers at the 75th percentile and workers with median earnings in that industry, and with a decrease in the gap between workers earning the median wages and those at the 25th percentile. This pattern can be explained by the tasks performed by workers. Offshoring is associated with a decrease in wages for occupations that rely heavily on routine tasks and an increase in wages if the occupation is nonroutine and communication task intensive. The results hold in both OLS and instrumental variable specifications.

*The statistical analysis of firm-level data on U.S. multinational companies was conducted at the Bureau of Economic Analysis, U.S. Department of Commerce under arrangements that maintain legal confidentiality requirements. The views expressed are those of the author and do not reflect official positions of the U.S. Department of Commerce. The author is grateful to seminar participants at the World Bank, Census, USITC, Trans-Pacific Labor Seminar, DOL Centennial Conference, and Georgetown University for helpful suggestions and to William Zeile and Raymond Mattaloni for assistance with the BEA data.

1 Introduction

The offshoring of intermediate inputs and back office services has received a great deal of attention in policy discussions and the popular press, especially as it relates to employment and wages in the US. Yet little empirical research has been done to increase our understanding of how offshoring impacts different types of workers and what the implications are for US labor markets. In this paper, I begin with a model proposed by Acemoglu and Autor (2010) that divides workers into high, medium, and low skilled groups. This model has been used to show how skill biased technology change can lead to a polarization of the workforce, that is, relative gains for workers at the high and low ends of the wage distribution and relative losses for workers in the middle. Acemoglu and Autor suggest that a similar result could be obtained from offshoring rather than technology change, however they do not empirically test that prediction. In this paper, I provide evidence that offshoring is indeed linked to increasing polarization of the U.S. workforce. I then show that the intensity with which occupations use nonroutine tasks determines which workers gain or lose from offshoring. To address the potentially endogenous relationship between US wages and offshoring by US firms, I instrument for US offshoring using a measure of offshoring by European firms. Firms in Europe are subject to the same exogenous technology shocks that have reduced the cost offshoring certain goods and services over time, yet their decisions are not directly dependant on wage fluctuations in the US.

Figures 1 and 2 illustrate the polarization that occurred in the U.S. labor force from 2001 to 2010. Figure 1 shows the log ratio of wages for workers at the 90th and 50th percentiles of the wage distribution, as well as the same ratio for workers at the 50th relative to the 10th percentiles. The gap between the 90th and 50th percentiles increased steadily over this time period, while the gap at the lower end of the distribution decreased slightly. Together, these two trends add up to relative losses for the median worker. Figure 2 shows the changes in log hourly earnings rates for the 10th, 25th, 50th, 75th, and 90th percentiles of the wage distribution from 2001 to 2010. The greatest gains have been at the top, with smaller gains at the bottom and middle of the distribution.

Previous work by Acemoglu and Autor (2010) and others has suggested that offshoring may contribute to this polarization trend, however that relationship has never been documented empirically. Traditional models of trade do not offer much insight into any relationship that may exist between offshoring and polarization. Increases in the gap between workers at the very top and very bottom of the wage distribution can be linked to offshoring

using standard comparative advantage models of trade. However, models of comparative advantage based on endowments of skill, capital, and other factors can not predict the pattern of polarization that we observe in the data. In this paper I present an empirical framework that reconciles this disconnect between theory and evidence. When comparative advantage is viewed in terms of tasks as well as labor endowments, more nuanced patterns in the data, such as polarization, can be explained.

Implicit in the speculations of previous authors that offshoring may lead to a polarization of labor demand in the US are a number of assumptions about which types of industries, occupations, or tasks can be more easily offshored. Fragmenting the production process in a way that moves certain intermediate inputs or processes to another country is easier for certain types of production than for others. For example, activities that involve direct interaction with consumers, such as haircuts or food service, are extremely difficult to offshore while those that are not location specific, such as data entry, can be performed anywhere. In addition, the literature¹ has shown that firms are more likely to offshore routine tasks and keep nonroutine tasks in their headquarters. Routine tasks can easily be broken down into a clear set of steps, which can then be programmed into computer code (in the case of skill biased technology change) or communicated to someone located in another country (in the case of offshoring). Nonroutine tasks involve decision making, problem solving, or creativity. These nonroutine tasks are both more difficult to communicate to overseas contractors or foreign affiliates and more crucial to a firm's core mission or strategy, and thus they are more likely to be performed in a company's US headquarters rather than offshored.

Figure 3 shows the importance of face-to-face communication for US workers by wage quintile, using wage data by occupation from 2002 and communication intensity scores by occupation from the Department of Labor's Occupational Information Network (O*NET). These measures follow a pattern that is very similar to that of the wage growth numbers used to document the polarization of labor demand. The most communication intensive industries are at the top and bottom of the wage distribution, with the least interactive, and therefore potentially the most offshorable, occupations concentrated in the middle of the distribution. Figure 4 shows the importance of nonroutine tasks for occupations divided by their 2002 wage quintile. The most nonroutine occupations are concentrated at the top of the wage distribution. However, we observe a u-shape at the lower wages levels, as the lowest quintile occupations are more intensive in nonroutine tasks than the second lowest quintile.

¹See Oldenski (2013) or Antras, Garicano and Rossi-Hansberg (2006)

But these tradeability considerations are not the only factors that determine which tasks will be offshored. At the same time that firms are making organizational decisions related to the ease of separating production tasks across borders, they also face the comparative advantage reality that workers in certain countries are relatively more productive at performing certain activities. When firms increase production at their foreign affiliates, they are not doing so in a vacuum. Table 1 shows that increases in employment at, or sales by, the foreign affiliates of US firms are associated with increases in both employment by those same firms in their US locations and with US exports by those firms. A 10 percent increase in employment at the foreign affiliates of a US multinational is associated with a 4 percent increase in both employment in the US headquarters of those firms and a 4 percent increase their exports from the US. The results presented in Table 1 are simply correlations and are not meant to imply a causal relationship, as it is difficult to disentangle the endogenous relationship between parent and affiliate activity. It is likely that decisions to increase activity at both foreign affiliates and in the US are made jointly. However, there could also be a causal link to the extent that increasing foreign affiliate activity increases productivity in a way that leads to higher employment in the US. The causality could also move in the opposite direction if exports and greater US employment follow foreign affiliate activity, for example, to provide post sale services or parts and maintenance. Whatever the mechanism, the evidence is strong that foreign affiliate and US activity within a firm move together and are compliments rather than substitutes. Any model that examines the loss of certain US jobs as a result of offshoring without considering the concurrent gains in other areas neglects this empirical fact.

Interactions between the “offshorability” considerations described above and comparative advantage can explain both the overall widening of the wage gap in the US, as well as disproportionate gains for workers at the top of the skill distribution and disproportionate relative losses to workers in the middle of the distribution. Nonroutine and communication intensive occupations are concentrated at both the top and the bottom of the skill distribution in the US. However, US workers at the top of the distribution also perform the jobs in which the US has comparative advantage. Thus we should expect these high skilled, high wage workers to gain the most from the increased tradeability of certain tasks both because the areas they specialize in are activities that are more in demand in the US as trade expands and because they are made more productive by the availability of cheaper complementary low skilled tasks through offshoring. Middle skilled workers should gain the least, as they are not in occupations for which the US has comparative advantage and are also more likely to

perform routine, easily offshorable tasks. Workers at the bottom of the distribution are not in sectors in which the US has comparative advantage, yet many of the tasks they perform are either nonroutine manual or location specific and thus less vulnerable to offshoring than middle skill jobs.

In this paper, I first document that a significant relationship exists between offshoring by US firms and the polarization of wages in the US from 2002 to 2008. Because the offshoring decisions of US firms may be endogenously related to US wages, I also instrument for changes in US offshoring using changes in offshoring by European countries. In doing so, I am attempting to isolate exogenous shocks to the offshorability of certain industries, such as technology changes that would impact firms anywhere in the world. When European firms are deciding whether and how much to offshore, US wages are not likely to be a major factor in this decision making process. Yet exogenous shocks that make offshoring relatively easier in some industries should impact firms from all countries. For example, information technologies and global supply chain management techniques expanded rapidly over the sample period. These technologies are available to both US and European firms. Thus if, say, internet technologies allow for firms to offshore back office services that were not previously offshorable, then this development should impact firms in both Europe and the US. If China's entry into the WTO is responsible for an increase in FDI in China by US firms, then the same should be true of European firms. Instrumenting for US firm behavior using the behavior of firms in other countries is an approach that has been used in the trade policy literature by Broda, Limao and Weinstein (2008) and Ludema and Mayda (2011). However, those papers are concerned with the impact of industry concentration on tariffs, rather than on offshoring and wages, and they instrument for the concentration of industries in a given country using the average concentration of each industry in similar countries, rather than using offshoring by European firms to instrument for US offshoring.

I then show that the tasks performed by workers in different occupations can explain the relationship between offshoring and polarization. The impact of offshoring on wages differs for occupations that are more or less routine or communication intensive and rely to a greater or lesser extent on the use of computers, even when controlling for skill and tradability, in a way that is consistent with an explanation based on an interaction between offshorability and comparative advantage. The results show that workers performing nonroutine tasks have experienced increases in their wages and employment shares independent of offshoring. However, it is these nonroutine occupations that have also gained the most from offshoring. While an increase in offshoring in a given industry reduces wages for the most routine occu-

pations in that industry, it increases wages and employment shares for the more nonroutine occupations. This suggests that U.S. workers have a comparative advantage in performing nonroutine tasks and that demand for workers performing these tasks increases in the U.S. when U.S. firms offshore more routine tasks.

2 Literature Review

Theoretical literature on the relationship between trade and wages dates back to the Stolper-Samuelson theorem of the Heckscher-Ohlin Model of trade. Simply put, in a two-country, two-good, two-factor world, an expansion of trade should lead to an increase in the returns to a country's abundant factor and a decrease in the returns to its scarce factor. Empirical tests of this prediction have produced mixed results.²

In the 1990's, the focus shifted to the impact of international outsourcing, as opposed to trade in final goods, on the wage gap³. However, these studies focused exclusively on the offshoring of manufactured goods (rather than services) and only divided workers into two skill groups based on educational attainment or their status as production versus nonproduction workers. Based on this classification, it has been well documented that the wage gap between high and low skilled workers has been growing over the last several decades and that international trade in intermediate goods has been among the contributing factors to this growing divide.

Over the past decade or so, offshoring has increasingly expanded beyond manufactured goods to include business and IT services. Amiti and Wei (2005) and Gorg and Hanley (2005) find that there are no long term or aggregate effect of an increase of service imports, but when highly disaggregate data are used, workers in some narrowly defined industries are negatively impacted in the short run. More recently, a few studies of trade and wages have begun to estimate the impact of trade and offshoring on incomes in the U.S. at a more detailed level. Ebenstein et al. (2011) use data from the Current Population Survey (CPS) and show that the impact of offshoring on wages and employment can be either positive or negative depending on the country in which the offshored work takes place and whether workers remain in the same sector or shift from manufacturing to services after the offshoring takes place. Liu and Trefler (2008) perform a similar exercise in which they examine the impact of an increase in service imports from India and China on employment, earnings, and

²See Leamer (2000) for a summary of empirical tests of the Stolper-Samuelson theorem

³See, for example, Feenstra and Hanson (1996, 1997 and 1999)

occupational switching using CPS data. However, these studies do not look at the overall distribution of wages.

On the theoretical side, an emerging literature has focused on trade in tasks, rather than in physical goods. These models are meant to capture the increasingly complex divisions of labor across countries that have been made possible by innovations in information technology. Several theoretical papers have been written to model the impact of this trend on wages.⁴ However, these papers do not address polarization and they have not been tested empirically, due primarily to the difficulty of acquiring data on trade in intangible services.

It is extremely difficult to test models of trade in tasks empirically because trade data are collected at the industry and firm level, rather than at the task or occupational level. Oldenski (2012 and 2013) has gotten around this data problem by constructing measures of the intensity with which industries use certain types of tasks. This work has shown that US firms are more likely to offshore routine tasks, while more nonroutine activities are performed in their US headquarters. Blinder (2009) has taken a more subjective approach, arguing that the most important factor in determining the offshorability of an occupation is the extent to which that occupation requires face-to-face interaction between the producer and the consumer. Jensen and Kletzer (2007) highlight the enduring importance of comparative advantage, offering evidence that even when certain activities are tradeable, that does not necessarily imply that they will be offshored by US firms, as the US is a net exporter of most tradable services. Firpo et al. (2011) show that the task content of occupations can explain changes in the US wage distribution. However, they do not use data on offshoring, but rather speculate that variation in the offshorability of tasks may explain this relationship. Thus they document a link between tasks and changes in the wage distribution, but do not empirically test for a relationship between tasks and offshoring or offshoring and polarization.

The labor literature has also devoted much attention to the rising wage gap between skilled and unskilled workers in the U.S. Studies by Berman, Bound and Griliches (1994), Autor, Katz and Kruegar (1998), and others have used what Acemoglu and Autor (2010) refer to as the canonical model to examine the impact of technology change on the wage gap. Recent work by labor economists has gone beyond the simple division of high and low skilled labor to investigate more subtle changes in the U.S. wage structure. Papers by Autor, Katz and Kearney (2006), Goos, Manning and Salomons (2009), Acemoglu and Autor (2010), and Autor and Dorn (2010) have documented a “missing middle” or polarization of the U.S. labor force. This polarization has been shown to be related to skill biased technological change.

⁴See, for example, Grossman and Rossi-Hansberg (2008) or Antras, Garicano and Rossi-Hansberg (2006)

Over time, returns to highly educated workers performing nonroutine cognitive tasks have been increasing, as those workers are complements to, and thus made more productive by, technological innovations. This leads to an increase at the top of the wage distribution. Workers performing routine tasks have seen their wages fall as many of the activities they perform are substitutes for technology. However, these workers are not concentrated at the very bottom of the wage distribution. Rather they represent many higher end manufacturing jobs as well as more routine service jobs, such as bank tellers and data entry workers. Jobs at the lowest end of the wage distribution include many that require nonroutine manual tasks, such as janitorial services and truck drivers, that can not be fully automated and replaced by machines. Thus these jobs have not experienced a fall in wages relative to those in the middle of the wage distribution. Offshoring has been suggested as an additional potential explanation for this polarization effect, but has not been included in empirical studies of the trend, which focus primarily on skill biased technology.

This paper attempts to bridge the gap between the approaches taken by labor economists and those taken by trade economists by combining detailed firm level data on the offshoring activities of U.S. multinationals with occupational level data on employment, wages, tasks and education.

3 Theoretical Framework

The basic framework follows Acemoglu and Autor (2010). They present extensive evidence of the recent polarization of the U.S. labor market, in which median wages have fallen relative to those of workers at both the top and the bottom of the wage distribution. They then propose a model with three types of labor and a continuum of tasks. Workers are either low, medium, or high skilled. Each skill level has comparative advantage in the performance of a subset of production tasks. Changes in technology (or in this case offshoring) may impact the productivity of one or more skill group in performing each task. Their focus is on the impact of skill-biased technical change, however they mention that a similar result can be obtained using offshoring rather than technology as the source of the productivity change. I begin by presenting their basic model, which I will test in the context of offshoring in Section 6.

3.1 The Acemoglu and Autor Model

Consider an economy with one final good Y that is produced by combining a continuum of tasks, $i \in [0, 1]$ using Cobb-Douglas technology

$$Y = \exp \left[\int_0^1 \ln y(i) di \right] \quad (1)$$

Each task has the following production function:

$$y(i) = A_L \alpha_L(i) l(i) + A_M \alpha_M(i) m(i) + A_H \alpha_H(i) h(i) \quad (2)$$

where L, M, and H index low, medium and high skilled workers. A is factor-augmenting technology and α is the productivity of workers at each skill level in performing task i . The linear structure of the task production function means that workers of any skill level can perform any task, however only one skill level will actually be used in production of that task in equilibrium. In this sense, it is a comparative advantage model of task performance. Tasks are indexed such that high skilled workers have comparative advantage in the higher numbered tasks and $\alpha_L(i)/\alpha_M(i)$ and $\alpha_M(i)/\alpha_H(i)$ are assumed to be continuously differentiable and strictly increasing. Assuming labor market clearing, it is straightforward to show that there will be two cut off points, I_L and I_H such that all tasks with indices $i < I_L$ will be performed by low skilled workers, all tasks with $i > I_H$ will be performed by high skilled workers and $I_L < i < I_H$ will be performed by medium skilled workers. Using the assumption of perfect competition in labor markets and properties of the Cobb-Douglas production function, relative wages across skill groups can be expressed as a function of the equilibrium allocation of tasks across skill groups, given by I_L and I_H :

$$\frac{w_H}{w_M} = \left(\frac{1 - I_H}{I_H - I_L} \right) \left(\frac{H}{M} \right)^{-1} \quad (3)$$

and

$$\frac{w_M}{w_L} = \left(\frac{I_H - I_L}{I_L} \right) \left(\frac{M}{L} \right)^{-1} \quad (4)$$

where H, M and L are number of high, medium and low skilled workers. The cutoff points I_L and I_H are in turn impacted by factor augmenting technology. Because of the comparative advantage nature of the model

$$\frac{dI_H}{d\ln A_H} < 0, \frac{dI_L}{d\ln A_H} < 0, \text{ and } \frac{d(I_H - I_L)}{d\ln A_H} < 0 \quad (5)$$

In other words, a positive shock to the productivity of high skilled workers increases the range of tasks performed by these workers, primarily at the expense of middle skilled workers. Thus the model predicts polarization as a result of skill biased technology change (i.e. innovations that increase the productivity of high skilled workers, A_H).

In an extension to this model, Acemoglu and Autor consider what will happen to relative wages if a share of the tasks used in production can be offshored to a foreign country. They assume that a set of tasks $[I', I''] \subset [I_L, I_H]$ can now be offshored. This assumption is based on the observation that an increasing number of relatively routine information tasks (such as data entry, basic accounting, and customer service) are being offshored while most extremely low skilled jobs (such as fast food workers and janitorial services) and extremely high skilled jobs (such as managers) are less likely to be offshored. After offshoring tasks in the range $[I', I'']$, new thresholds \hat{I}_L and \hat{I}_H will exist such that $\hat{I}_L < I_L$ and $\hat{I}_H > I_H$. This will result in an increase in w_H/w_M and a decrease in w_M/w_L . However, they do not empirically test this result. This is the first prediction I will be testing in Section 6.

In addition to the argument about the relative ease of offshoring certain types of tasks that are more often performed by middle skilled workers, there is also a link between traditional skill-based comparative advantage and polarization. Table 1 presents evidence that offshoring does not take the form of a simple net gain in foreign jobs at the expense of a net loss in domestic jobs. As some work moves offshore, there are other areas of employment in the US that expand. As a country increases its trade for any reason, be it free trade agreements, falling transport costs, or innovations that allow for greater fragmentation of production, the areas in which the US has comparative advantage will gain the most. In the US, a skill abundant country, high skilled workers benefit. Under the framework described above, in which there are three types of workers, a continuum of tasks, and an equilibrium in which each task is performed by only one worker, an increase in trade (or offshoring) will result in polarization. An increase in the offshorability, or more accurately the tradeability, of tasks can be modeled as an increase in α_H relative to α_L and α_M due to the US comparative advantage in high-skill intensive products. Similar to the skill biased technical change result, we have polarization because

$$\frac{dI_H}{d\ln \alpha_H} < 0, \frac{dI_L}{d\ln \alpha_H} < 0, \text{ and } \frac{d(I_H - I_L)}{d\ln \alpha_H} < 0 \quad (6)$$

In this case, the change in α_H works through exports, as increased tradability leads to increased specialization in the areas where the US has comparative advantage.

However, if it is also the case that the increase in tradeability does not impact tasks at all skill levels equally, then the impact also comes through yet another channel, as shown by Acemoglu and Autor in the case in which only tasks $[I', I''] \subset [I_L, I_H]$ can be offshored. These two effects together (traditional comparative advantage based on skill endowments and variations in tradeability that are correlated with skill) can predict polarization of the US labor force as a result of offshoring, as well as the overall widening of the wage gap.

4 Empirical Specification

4.1 Industry Level

I begin by testing the prediction of Section 3.1 that an increase in offshoring will lead to an increase in gap between workers at the top of the wage distribution and workers in the middle of the wage distribution. At the same time, the wage gap between the median worker and workers at the bottom of the distribution will narrow. For the industry level exercise, I estimate

$$\Delta \ln \left(\frac{w_{75}}{w_{50}} \right)_{kt} = \alpha + \beta_1 \Delta \ln(\text{trade})_{kt} + \beta_2 \Delta \ln(\text{tech})_{kt} + \gamma_t + \varepsilon_{kt} \quad (7)$$

and

$$\Delta \ln \left(\frac{w_{50}}{w_{25}} \right)_{kt} = \alpha + \beta_1 \Delta \ln(\text{trade})_{kt} + \beta_2 \Delta \ln(\text{tech})_{kt} + \gamma_t + \varepsilon_{kt} \quad (8)$$

Where $(\frac{w_{75}}{w_{50}})_{kt}$ is the ratio of wages at the 75th percentile to wages at the median of the earnings distribution in industry k at time t and $(\frac{w_{50}}{w_{25}})_{kt}$ is the ratio of wages at the median to wages at the 25th percentile in industry k at time t . The variable *trade* refers to offshoring and exports. It is measured first by the share of total U.S. production in each industry that takes place at foreign affiliates of U.S. multinational firms. But as shown in Table 1, offshoring does not occur in a vacuum, as imports of some goods and services are complemented by exports of others. Thus the empirical specifications will also look at exports by US MNCs as well as offshoring to isolate the comparative advantage forces underlying the relationship between offshoring and polarization. *Tech* captures skill biased technical change. It is measured using the fraction of workers employed in each industry

who report using a computer at their job in the CPS 2003 October Supplement. This variable captures the technology intensity of the industry, which has been shown to impact the polarization of wages. I am able to construct a time-varying measure of computer use at the industry level by starting with the raw data at the occupation level and exploiting changes in occupational employment shares within industries over time when constructing the industry-level variable. γ_t is a time fixed effect. Time invariant industry-level characteristics are controlled for through the use of first differences.

4.2 Occupation Level

After testing for a link between offshoring and polarization, I will turn to more detailed task and occupation characteristics to explain this relationship. Offshoring will likely have different effects on different occupations, and may also impact the distribution of wages *within* an occupation depending on the skill level, routineness, communication intensity, and computer use required for that occupation. Technology that allows for the unbundling of tasks may not just involve offshoring the entire occupation of a single worker. It may also result in individual workers having some of the tasks they perform offshored, freeing them up to focus on higher value added tasks where they have greater comparative advantage and can be more productive. For example, an economist does research but may also type the results of that research. A management consultant may create some novel client-specific solutions, but likely also spends some time simply packaging and presenting publicly available information. The relative gains from offshoring should be highest for the most highly skilled workers in the US, as they have the greatest relative gains from shedding lower value added tasks when the US increases its specialization in nonroutine tasks.

Even within a single occupation, we can expect to see variations in who gains and who loses from offshoring. Workers in the same occupation vary in terms of their relative productivity at performing the various tasks required for that occupation. The workers within an occupation who are better at performing the most routine tasks will likely be hurt the most by offshoring, while those workers with comparative advantage in nonroutine tasks will gain the most. Occupations that require a larger share of nonroutine tasks have a greater scope for specialization according to comparative advantage as a result of trade, thus we should see greater within occupational shifts among occupations that are more intensive in their use of nonroutine tasks as comparative advantage across tasks within an occupation takes hold, as well as aggregate relative gains to those occupations compared to occupations that

specialize in routine tasks.

An increase in offshoring of some of the tasks associated with a given occupation will decrease the relative demand for workers with comparative advantage in those tasks, having a negative impact on their relative wages. This leads to distributional changes within, as well as across, occupations.

Workers who had previously performed both routine and nonroutine tasks can now allocate a greater share of their labor hours to the nonroutine tasks, magnifying the comparative advantage of workers who are most productive at performing those tasks.

To test for these effects of task offshoring within and across occupations, I estimate:

$$\begin{aligned} \Delta \ln(X)_{jkt} = & \alpha + \beta_1 \Delta \ln(\text{offshoring})_{kt} + \beta_2(\text{tech})_j + \beta_3(\text{task})_j \\ & + \beta_4 \Delta \ln(\text{offshoring})_{kt}(\text{tech})_j + \beta_5 \Delta \ln(\text{offshoring})_{kt}(\text{task})_j + \gamma_t + \gamma_k + \varepsilon_{jkt} \end{aligned} \quad (9)$$

Where X represents a variety of outcome variables of interest, including averages wages, employment, and polarization measures at the occupation-industry level. Because offshoring data are not available by occupations, I look at the impact of an increase in offshoring at the industry level on the outcomes of occupations employed in that industry. So, for example, X_{jkt} may be the average wage of administrative assistants employed in the legal industry, in which case offshoring_{kt} would be offshoring done by the legal industry. nonroutine is the importance of nonroutine tasks in the occupation. γ_t and γ_k are time and industry fixed effects.

4.3 Instrumental Variable Approach

It is possible that the offshoring decisions of US firms could be endogenous to US wages. This would be the case if US wages for certain occupations were to rise or fall enough to alter the relative costs of producing certain inputs at home versus abroad in a way that impacted offshoring decisions. Firms have been shown to consider many factors other than wages in their offshoring decisions, such as tax policy, institutions, technological feasibility, and trade agreements. But wages may also play a role. The impact of changes in US wages on offshoring is likely to be small, as the differences in wages across countries are much greater than the differences in wages in the US from one year to the next. However, it is difficult to rule out the possibility that the changes in US wages during the sample period may have had some impact on the offshoring decisions of US firms. Thus I instrument for US

offshoring using offshoring by European firms. When European firms are deciding whether to offshore, they are considering wages in their own home countries relative to wages in the rest of the world. This decision is not tied directly to US wages the way that offshoring by US firms may be. Yet exogenous shocks that make offshoring relatively easier in some industries should impact firms from all countries. For example, information technologies and global supply chain management techniques expanded rapidly over the sample period. These technologies are available to both US and European firms. Thus if, say, internet technologies allow for firms to offshore back office services that were not previously offshorable, then this development should impact firms in both Europe and the US. If China's entry into the WTO is responsible for an increase in offshoring to China by US firms, then the same should be true of European firms. Because the variation in the US offshoring data is at the industry level, I need a measure of European offshoring that also varies by industry. The OECD publishes data on FDI flows for member countries at both the country and industry level. Unfortunately, not all OECD countries report flows at the industry level. Thus I use the weighted percentage change in outward FDI flows for three countries that do consistently report FDI by industry: the United Kingdom, Germany, and France. These changes are correlated with changes in offshoring by US MNCs, yet they are not significantly correlated with US wages.

A similar approach has been used by Broda, Limao and Weinstein (2008) and Ludema and Mayda (2011). Both of those papers study terms of trade effects of tariffs, rather than the relationship between offshoring and US wages. However, they have similar endogeneity concerns when looking at the role of market power in one country on the tariffs set by that country. Both papers instrument for a given country's market power in a good using the average market power (or industry concentration) in that good in other similar countries.

5 Data

The Bureau of Economic Analysis collects firm-level data on U.S. multinational company operations in both goods-producing and service-producing industries in its annual surveys of U.S. direct investment abroad. I use these data to define a measure of offshoring. This variable consists of the total sales by a foreign affiliate of a U.S. multinational as a share of the firm's total sales. This variable captures the share of total output produced by U.S. owned firms that is produced in countries other than the U.S. Sales by foreign affiliates include sales back to the U.S. as well as to the country in which the affiliate is located and

the rest of the world. I use data from 2002 to 2008 because these are the years for which occupational employment and wage data are also available. The BEA surveys cover 112 manufacturing and service industries, classified according to BEA versions of 4-digit North American Industrial Classification System (NAICS) codes.

The instrumental variable is constructed using data on industry level outward FDI flows from Germany, France and the United Kingdom, which are collected by the OECD.

Data on employment and wages come from the Occupational Employment Statistics (OES) put out by the Bureau of Labor Statistics (BLS). The OES provides employment and wage data by occupation and industry from BLS payroll surveys. They also include information on wages for workers at the 10th, 25th, 50th, 75th, and 90th percentiles by both occupation and industry.

Information on computer usage comes from the Current Population Survey (CPS) October 2003 Supplement. The CPS provides individual level data on workers, including their industry and occupation of employment, wages, and education level. The October Supplement also asks whether the worker uses a computer or accesses the internet as part of his or her job. Unfortunately, the computer questions are not included in the October CPS Supplement every year, and the most recent year in which the question was asked was 2003. For the industry level analysis, I construct a time series variable that takes advantage of changes in occupational shares within each industry over time. I begin with the share of workers that used a computer at their jobs for each occupation in 2003. I then construct a separate industry-level variable for each year in the sample that weights these computer use shares by the share of workers in each occupation in the industry in each year. Thus an industry exhibits increasing computer intensity if it increases its share of workers in occupations that used computers intensively in 2003. For the occupation level analysis, I use the CPS data on computer by occupation use for 2003 for the entire sample period.

I also match the data on offshoring and employment with data from the Department of Labor's (DOL) Occupational Information Network (O*NET), which includes data on the importance of over 200 tasks or work activities performed in about 800 occupations. The information used by DOL to construct the O*NET measures is collected using a two-stage sampling design in which a statistically random sample of businesses expected to employ workers in the targeted occupations is selected and a random sample of workers in those occupations within those businesses is surveyed using standardized questionnaires.

To capture task routineness, I focus on the O*NET measure that most closely captures the importance of nonroutine tasks: the measure for "making decisions and solving problems."

To capture communication intensity, I use the measure for “interacting with the public.” For the occupation level analysis, I use the importance scores for each task, relative to the sum of the importance scores of all tasks used in the occupation. For the industry level analysis, I aggregate these scores up to the industry level weighting by the share of employment in each industry accounted for by each occupation in each year, resulting in a variable that varies by industry and year, as with the computer use measure described above. The importance score of each task, i in each industry, k is

$$M_{ik} = \sum_j \alpha_{kj} \ell_{ij} \quad (10)$$

where i indexes tasks, j indexes occupations, and k indexes industries. Thus α_{kj} is the share of occupation j used in the production of industry k , and ℓ_{ij} is an index of the importance of task i for occupation j .⁵ Summing over occupations in a given industry results in an index of the un-scaled importance score for each task in that industry. Each raw score is then divided by the sum of scores for each task in each industry, resulting in an input intensity measure for each task, i , in each industry, k :

$$I_{ik} = \frac{M_{ik}}{\sum_i M_{ik}} \quad (11)$$

Occupations are matched to industries using the Bureau of Labor Statistics Occupational Employment Statistics.

6 Results

6.1 OLS Industry Level Results

Before estimating the polarization equations (7) and (8), I first look at the impact of offshoring on average wages and employment shares by industry. Table 2 presents these results. A change in the log value of offshoring in a given industry has no impact on the average wage in that industry, either contemporaneously or with one or two year lags. I also control for the intensity with which each industry uses computers, a proxy for skill biased technical change.

⁵ ℓ_{ij} corresponds to the 0-100 score O*NET reports to measure the importance of each task in each occupation. These scores are constructed from surveys of individuals in those occupations and are normalized to a 0-100 scale by analysts at the Department of Labor. Due to the subjective nature of the surveys, one unit of importance for given task can not be directly compared to one unit of another task. This is a limitation of the data and motivates the use of relative intensity scores rather than the raw scores reported by O*NET.

Employment in the industry doing the offshoring as a share of total U.S. employment is not impacted by a change in offshoring in the current or previous year. However, two years after an increase in offshoring, the share of U.S. workers employed in the industry doing the offshoring increases. A change in the intensity with which computers are used in the industry has no effect on aggregate wages or employment shares, whether contemporaneously or with one or two year lags.

Even though there is no evidence that changes in offshoring impact the average wage paid to workers at the 4-digit NAICS industry level, they are associated with changes in the distribution of wages. Table 3 presents the results on polarization. Consistent with previous literature, I find evidence that increased computer use is associated with relative gains at the top and bottom of the earnings distribution and relative losses in the middle. Table 3 shows that the gap between the workers earning the median wage in an industry and the workers whose wages are at the 25th percentile for the industry decreases with an increase in computer use while the gap between wages at the median and the 75th percentile increases. The same relationship holds for offshoring and relative wages, however in the case of offshoring, the effect is only significant in the current period and with a one year lag. The effect is also stronger, more precisely estimated, and more persistent at the top rather than the bottom of the wage distribution. These results are consistent with evidence of polarization found in the labor literature. In other words, the occupations that are least in demand as a result of offshoring are those in the middle of the distribution, while demand increases for workers at the bottom and the top of the distribution, either because these tasks are more difficult to offshore, or because they are areas in which U.S. workers have a comparative advantage and are thus in higher demand when trade expands and productivity rises.

Because the regressions are run in logs, the coefficients can be interpreted as percentage changes. So, for example, a 10 percent increase in the share of activity in an industry that takes place at foreign affiliates is associated with a 0.3 percent increase in the gap between workers at the 75th and 50th percentiles of the wage distribution in that industry, and with a 0.9 percent decrease in the gap between workers at the 50th and 25th percentiles. In 2002, 32% of MNC activity took place at foreign affiliates. By 2008 that share was 42%, a 31% increase relative to 2002. Using the coefficient on offshoring from Column 1 of Table 3, this translates into about a 0.9% increase in the wage gap at the top of the distribution that can be explained by offshoring. During the sample period, the ratio of wages at the 75th relative to the 50th percentiles of the wage distribution in the US increased by 1.6%. Thus

offshoring can explain slightly more than half of the increase in the wage gap at the top of the distribution.

6.2 OLS Occupation Level Results

Having documented the relationship between offshoring and polarization above, I now turn to detailed explanations for why this relationship exists and what factors are driving the differential impact of offshoring on different types of workers. Specifically, I examine how the link between offshoring and polarization varies based on the skill level, computer usage, and task content of occupations.

When data are disaggregated to the six-digit occupation level, the link between offshoring and both average wages and employment is much stronger. Table 4 shows that when an industry exhibits a greater level of offshoring, the average occupation in that industry sees an increase in wages. There are contemporaneous falls in employment shares, but employment increases with one and two year lags. This is perhaps surprising if one thinks of offshoring as a static shifting of jobs from the US to another country. However, the results suggest that this static, one-direction view is not accurate. The wage effect may in part contain a composition effect if low wage jobs are being eliminated, leaving relatively more higher wages jobs to raise the average. But this effect should not be a concern for employment, which increases for occupations within an industry that exhibits offshoring growth. In Section 6.4 I control for the wage composition effect by defining each occupation's position in the wage distribution as constant using values for a base year and then looking at changes over time for occupations at different initial positions in the distribution. The results still hold.

These averages mask important differences across occupations. I now turn to measures of potential sources of comparative advantage and tradeability to understand what is driving the relationship between polarization and offshoring. Table 5 shows the effects of offshoring on occupations interacted with the computer use, task characteristics, and skill levels of those occupations. When interactions between computer use or tasks and offshoring are included, the coefficient value on a change in the offshoring share is negative. Workers in more computer intensive occupations are negatively impacted by offshoring in their industry. The results interacting offshoring and communication tell a similar story. If the offshoring share of an industry increases, workers in the more communication intensive occupations in that industry gain. These results are consistent with a tradability story. Since more communication-intensive jobs are more difficult to offshore, workers in these jobs gain from

offshoring while workers in less communication intensive occupations experience a fall in average wages.

However, comparative advantage also plays a role in addition to tradeability. Table 5 includes an interaction between a change in the offshoring share and the educational level of workers in each occupation. Consistent with a US comparative advantage in more highly skilled professions, the coefficient on the interaction between offshoring and education is positive, suggesting that workers in high skilled occupations experience wage benefits from offshoring, while workers in less skill-intensive occupations experience wage losses. Similarly, offshoring has a positive effect on the average wages of workers in more nonroutine task-intensive occupations. All of these results together can explain the broad patterns of increasing inequality and polarization that result from offshoring. High skilled and high wage workers gain from offshoring due to the US comparative advantage in the occupations they perform. Workers in less communication intensive and more routine occupations are made worse off. However, because these workers are concentrated in the middle of the wage distribution, the relationship between offshoring and average wages is not monotonically increasing with a worker's position in the wage distribution.

6.3 Controlling for Composition Effects of Offshoring

The above results use data on what the employment, average wages, and wages at various percentiles of the distribution were in each year. However, changes in these outcomes can come about in a number of different ways. For example, the average wage in an industry could increase because there is an overall increase in demand for workers in that industry. However, if the firms in an industry simply eliminated their lowest skilled, lowest paid workers, then the average wage would also increase, but in this case due to a composition effect rather than an aggregate demand increase. Similarly, the wages at each percentile of the distribution may be endogenous to offshoring. A given occupation that was at, say, the 25th percentile of the wage distribution in one year may not occupy that same place in the distribution a few years later if other occupations have been eliminated or expanded. It is still important to look at the relationship between offshoring at relative wages at any given point in time, as the wage gap in any given year is still meaningful regardless of whether it came about through exogenous or endogenous shifts. However, we can have an even better understanding of this relationship by also looking at changes in wages relative to the distribution in some fixed base year, instead of allowing the ranking of occupations within the distribution to

change as well.

To address these issues, I reran the primary specifications defining each occupation's place in the wage distribution using data from 2002. I assigned each occupation to either the lowest, middle, or highest third of the wage distribution based on average earnings in 2002. I then looked at how changes in offshoring and computer use have impacted these three groups differently over time. Table 6 shows how average wages changed over time for workers who were in the lowest, middle, and highest wage occupations based on their 2002 earnings. The results in columns 1-3 show that offshoring had a small positive impact on workers in the lowest wage industries. Workers in the middle third were negatively impacted as offshoring increased. And workers in the top third experienced wage gains in association with offshoring. The results are qualitatively similar to those for changes in computer usage by industry over time.

The extent to which each occupation relies on nonroutine tasks plays a key role in the relationship between offshoring and wage changes. As shown in columns 4-6 of Table 6, while offshoring has an overall negative impact on workers in occupations that are in the bottom third of the wage distribution, it has a positive impact on wages in the more communication intensive low wage occupations. Nonroutine intensity is not significant for the middle third of industries, possibly due to the small size of the sample and the relatively low intensity of nonroutine tasks for middle income jobs. For workers in the highest paying industries, offshoring is associated with wage gains for the most communication intensive, highly educated, and nonroutine occupations.

6.4 Distributional Effects of Offshoring Within Occupations

All of the above results, both at the industry and the occupational level, capture shifts across different types of jobs. At the industry level, offshoring does not have a significant aggregate effect, but it does increase demand for the highest paid workers in an industry, while decreasing relative demand for workers at the middle of the distribution. At the occupation level, offshoring industries increase demand for workers in nonroutine task intensive occupations but decrease demand for workers in more routine task intensive industries. However, offshoring does not necessarily take place at the occupational level. In other words, as a company expands its foreign operations, it is not necessarily going to have workers at the foreign affiliate perform identical tasks to those that were performed by workers in the US headquarters. Consider an example from the accounting sector. An accountant could per-

form any of a wide range of accounting tasks, from simply filling in and filing tax returns to providing long term planning advice. Suppose that an accounting firm has a number of employees who each perform a range of these tasks. Now suppose that in order to cut costs, this accounting firm offshores some of its work to an affiliate in another country, say, India. It could do this by hiring employees in India to do exactly what the employees in the US had been doing. Or it could have the employees in India perform a subset of the tasks previously performed in the US, leaving the accountants in the US to focus on the remaining tasks. Anecdotal and case study evidence suggests that this is indeed the way many firms implement offshoring (Friedman 2005). Previous empirical work on offshoring by Oldenski (2012) shows that US firms offshore the most routine tasks, keeping more nonroutine activities in their headquarters locations.

All of this evidence suggests that the appropriate units to use when measuring offshoring should be tasks, rather than occupations or industries. Unfortunately, trade, employment, and wage data are not collected at the task level. However, looking at how offshoring impacts the distribution of wages among workers in the same occupation can provide evidence of within-occupation shifts that are likely correlated with the types of tasks performed by different workers within those occupations. Consider the accounting example described above. In this case, offshoring is likely to have a negative impact on accountants whose specialize in routine tasks. Accountants with comparative advantage in nonroutine tasks should then comprise a larger share of demand for accountants in the US. To the extent that offshoring increases the productivity of the firm doing the offshoring, allowing them to expand their operations in all locations, this could result in net gains for nonroutine task performing accountants in the US, at the expense of US accountants who have a comparative advantage in routine tasks.

Table 7 captures these distributional effects within occupations. The results presented in Table 7 show that offshoring increases the wage gap of workers within the same occupation for occupations that are more nonroutine task intensive and that have higher educational requirements. This is true at both the bottom and the top of the distribution.

6.5 Instrumental Variable Results

Table 8 shows the first stage results using offshoring by European firms to instrument of US offshoring. As expected, industry level changes in the FDI volumes of European MNCs are positively correlated with changes in the activities of US MNCs at their foreign affiliates.

Table 9 repeats the industry-level polarization exercise using the European FDI instrument. The results still hold. An increase in the offshoring share of a given industry leads to an increase in the gap between wages paid to workers at the 75th percentile and the 50th percentile within that industry. It also leads to a reduction in the gap between wages paid to workers at the 50th and 25th percentiles. These two results combined show that offshoring is associated with the increasing polarization of the US labor market.

Table 10 shows how this dynamic plays out across occupations within a given industry. As with the OLS specification, when task requirements are controlled for, an increase in offshoring in an industry leads to a decrease in the average wages of occupations in that industry. However, task intensities of the occupation also matter. Offshoring has a positive impact on the wages of occupations that are high skilled and intensive in nonroutine or communication tasks. Because these tasks are concentrated at the top and bottom of the wage distribution, the net effect is a link between offshoring and polarization at the occupation level. Offshoring leads to a reduction in wages for middle skilled occupations and an increase in wages for high skilled occupations, with no significant change at the bottom of the distribution.

Table 11 shows that the results hold at the occupation level when both 2SLS and the division on occupations based on their position in the 2002 wage distribution are used.

6.6 The Role of Exports

As described in Section 3, there are two main channels through which offshoring impacts polarization. The first is related to variations in the technological feasibility of offshoring tasks at different positions in the wage distribution. As suggested by Acemoglu and Autor (2010), the most easily offshored tasks include manufacturing production and routine office work which are concentrated in the middle of the wage distribution. The lowest end of the wage distribution is mostly made up of nontradable occupations such as food service, janitorial work, and child care. Yet it is also true that the US is relatively well endowed with skilled labor, and thus has a comparative advantage in the occupations at the high end of the wage distribution. Thus as tradability expands, we should expect to see more highly skilled, highly paid workers gaining as a result of increased specialization in the tasks that they perform.

Table 12 reports the second stage industry level polarization results controlling for both offshoring and exporting. When exports are included the coefficient on offshoring is lower in

magnitude, suggesting that some of the impact of offshoring can be explained by increasing specialization along the lines of comparative advantage. However, the coefficient on offshoring is still significant, suggesting that offshoring plays a role beyond simply shifting away from importing industries and toward exporting industries.

Table 13 reports similar results using occupations divided into thirds based on their position in the 2002 wage distribution. As expected, increased exporting in an industry has a positive impact on the wages of the most highly skilled occupations within that industry. The results for offshoring still hold but are smaller in magnitude when offshoring is controlled for.

7 Conclusions

This paper shows that increases in offshoring at the industry level are associated with patterns of polarization in the U.S. labor market. This polarization can be explained by the importance of communication and nonroutine tasks, which are present at the top and bottom of the skill distribution to a greater extent than the middle of the distribution. In addition, offshoring leads to occupation-level gains in terms of average wages and employment. However, these gains are accompanied by increases in the wage gap for different workers, as those with comparative advantage in communication intensive and nonroutine tasks see greater returns to their skills, while those with comparative advantage in more routine tasks do not.

8 References

Acemoglu, Daron and David H. Autor. 2010. Skills, Tasks and Technologies: Implications for Employment and Earnings. Handbook of Labor Economics Volume 4, Orley Ashenfelter and David E. Card (eds.), Amsterdam: Elsevier.

Amiti, Mary and Shang-Jin Wei, 2005. Fear of service outsourcing: is it justified? Economic Policy, CEPR, CES, MSH, vol. 20(42), pages 308-347, 04.

Antras, Pol, Luis Garicano, and Esteban Rossi-Hansberg, 2006, Offshoring in a Knowledge Economy. Quarterly Journal of Economics 121(1), pp. 31-77.

Autor, David H., Lawrence F. Katz and Melissa S. Kearney. 2008. Trends in U.S. Wage Inequality: Revising the Revisionists. *Review of Economics and Statistics*, 90: 300-23.

Autor, David H., Lawrence F. Katz and Melissa S. Kearney. 2006. Measuring and Interpreting Trends in Economic Inequality: The Polarization of the U.S. Labor Market. AEA Papers and Proceedings.

Autor, David H., Lawrence F. Katz and Alan B. Krueger. 1998. Computing Inequality: Have Computers Changed the Labor Market? *Quarterly Journal of Economics*, 113: 1169-1213.

Autor, D, F. Levy, and R. Murnane, 2003, The Skill Content of Recent Technological Change: an Empirical Exploration. *Quarterly Journal of Economics* 118(4)

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. *The Quarterly Journal of Economics* 109(2), pp. 367-397

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

Blinder, Alan. 2006. "Offshoring: The Next Industrial Revolution?" *Foreign Affairs*

Broda, Christian, Nuno Limao, and David Weinstein. 2008. Optimal Tariffs and Market Power: The Evidence. *American Economic Review* 98(5). pp. 2032-2065.

Costinot, Arnaud, Lindsay Oldenski and James Rauch, 2011, "Adaptation and the Boundary of Multinational Firms," *Review of Economics and Statistics*.

Desai, M, F. Foley and J. R. Hines Jr., 2001, Repatriation Taxes and Dividend Distortions, *National Tax Journal*, 54, pp. 829-851.

Ebenstein, Avraham Y., Harrison, Ann E., McMillan, Margaret S. and Phillips, Shannon, 2009, Estimating the Impact of Trade and Offshoring on American Workers Using the Cur-

rent Population Surveys. NBER Working Paper w15107

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux, 2011, Occupational Tasks and Changes in the Wage Structure. Working Paper.

Freeman, Richard B. and Remco Oostendorp, 2000, Wages Around the World: Pay Across Occupations and Countries. NBER Working Paper No. W8058.

Friedman, Thomas, 2005, *The World is Flat: A Brief History of the Twenty-first Century*. Farrar, Straus and Giroux.

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

Gorg, Holger and Hanley, Aoife, 2005. "Labour demand effects of international outsourcing: Evidence from plant-level data," *International Review of Economics and Finance*, Elsevier, vol. 14(3), pages 365-376.

Grossman, G. and E. Rossi-Hansberg, 2008. "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98(5).

Jensen, J. B. and L. Kletzer, 2007, Measuring Tradable Services and the Task Content of Offshorable Services Jobs. In K. Abraham, M. Harper and J. Spletzer, eds., *Labor in the New Economy*, University of Chicago Press, forthcoming.

Kemeny, Thomas and David Rigby, 2010, Trading Away What Kind of Jobs? Globalization, Trade and Tasks in the U.S. Economy. Working Paper.

Leamer, Edward. 2000. What's the use of factor contents? *Journal of International Economics* 50(1), 17-49.

Liu, Runjua and Daniel Treffer. 2008. Much Ado About Nothing: American Jobs and the Rise of Service Outsourcing to China and India. NBER Working Paper 14061.

Ludema, Rodney and Anna Maria Mayda. 2011. Do terms of trade effects matter for trade agreements? Theory and evidence from WTO countries. Working Paper.

Melitz, M.J. 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71.

Oldenski, Lindsay, 2012, Export Versus FDI and the Communication of Complex Information. *Journal of International Economics*, 87(2), pages 312-322

Oldenski, Lindsay, 2013, The Task Composition of Offshoring by U.S. Multinationals. *International Economics* (forthcoming)

Table 1: Offshoring and onshore activity, 1990-2008

Model :	1	2	3	4
N:	41347	40931	21094	20592
Depvar:	$\Delta\ln(\text{par emp})$	$\Delta\ln(\text{par emp})$	$\Delta\ln(\text{par exp})$	$\Delta\ln(\text{par exp})$
$\Delta\ln(\text{affiliate emp})$	0.386*** (0.013)		0.396*** (0.029)	
$\Delta\ln(\text{affiliate sales})$		0.216*** (0.011)		0.324*** (0.028)
Year FE	yes	yes	yes	yes
R-sq	0.17	0.10	0.07	0.06

*, ** and *** indicate significance at the 10, 5 and 1 percent levels, respectively
Robust standard errors are in parentheses

Table 2: Industry level regressions: average wages and employment shares

Model :	1	2	3	4	5	6
N:	672	560	448	672	560	448
Depvar:	$\Delta \ln(\text{wage})$	$\Delta \ln(\text{wage})$	$\Delta \ln(\text{wage})$	$\Delta \ln(\text{emp})$	$\Delta \ln(\text{emp})$	$\Delta \ln(\text{emp})$
$\Delta \ln(\text{offshore})$	-0.078 (0.057)			0.061 (0.076)		
$\Delta \text{ comp}$	0.001 (0.002)			-0.001 (0.001)		
$t - 1 \Delta \ln(\text{offshore})$		-0.030 (0.062)			0.024 (0.079)	
$t - 1 \Delta \text{ comp}$		0.019 (0.030)			-0.038 (0.039)	
$t - 2 \Delta \ln(\text{offshore})$			0.018 (0.061)			0.134* (0.070)
$t - 2 \Delta \text{ comp}$			0.043 (0.040)			-0.061 (0.064)
R-sq	0.13	0.12	0.18	0.09	0.08	0.18

*, ** and *** indicate significance at the 10, 5 and 1 percent levels, respectively

Robust standard errors are in parentheses

Table 3: Industry level regressions: polarization

Model :	1	2	3	4	5	6
N:	672	560	448	672	560	448
Depvar:	$\Delta\ln(75/50)$	$\Delta\ln(50/25)$	$\Delta\ln(75/50)$	$\Delta\ln(50/25)$	$\Delta\ln(75/50)$	$\Delta\ln(50/25)$
$\Delta\ln(\text{offshore})$	0.026* (0.014)	-0.086** (0.038)				
Δcomp	0.041* (0.022)	-0.036* (0.019)				
$t - 1\Delta\ln(\text{offshore})$			0.104** (0.042)	-0.007 (0.042)		
$t - 1\Delta\text{comp}$			0.100** (0.031)	-0.021*** (0.002)		
$t - 2\Delta\ln(\text{offshore})$					0.019 (0.025)	0.003 (0.009)
$t - 2\Delta\text{comp}$					0.137*** (0.052)	-0.219*** (0.044)
R-sq	0.04	0.08	0.03	0.09	0.04	0.11

*, ** and *** indicate significance at the 10, 5 and 1 percent levels, respectively
Robust standard errors are in parentheses

Table 4: 6-digit occupation-level regressions: average wages and employment shares

Model :	1	2	3	4	5	6
N:	84749	77053	68855	77740	70827	63437
Depvar:	$\Delta\ln(\text{wage})$	$\Delta\ln(\text{wage})$	$\Delta\ln(\text{wage})$	$\Delta\ln(\text{emp})$	$\Delta\ln(\text{emp})$	$\Delta\ln(\text{emp})$
$\Delta\ln(\text{offshoring})$	0.002*** (0.001)			-0.009* (0.005)		
t-1 $\Delta\ln(\text{offshr})$		0.004*** 0.000			0.021*** (0.005)	
t-2 $\Delta\ln(\text{offshr})$			0.004** 0.000			0.012*** (0.004)
comp	-0.023 (0.032)	0.113*** (0.027)	0.082*** (0.028)	0.057 (0.307)	-0.674** (0.268)	-0.538** (0.221)
R-sq	0.13	0.13	0.13	0.14	0.15	0.13

*, ** and *** indicate significance at the 10, 5 and 1 percent levels, respectively

Robust standard errors are in parentheses

Table 5: Occupation level regressions

Model :	1	2
N:	68855	77740
Depvar:	$\Delta\ln(\text{wage})$	$\Delta\ln(\text{emp})$
t-1 $\Delta\ln(\text{offshoring})$	-0.010*** (0.002)	-0.050*** (0.018)
t-1 $\Delta\ln(\text{ofshr})^*\text{nrtne}$	0.081* (0.047)	0.548 (0.580)
t-1 $\Delta\ln(\text{ofshr})^*\text{comp}$	-0.059* (0.034)	-1.234*** (0.342)
t-1 $\Delta\ln(\text{ofshr})^*\text{comm}$	0.086** (0.035)	1.935*** (0.357)
t-1 $\Delta\ln(\text{ofshr})^*\text{edu}$	0.002*** (0.0004)	0.019*** (0.003)
R-sq	0.26	0.27

*, ** and *** indicate significance at the 10, 5 and 1 percent levels. Robust standard errors are in parentheses

Table 6: Occupation level regressions by position in the 2002 wage distribution

Model :	1	2	3	4	5	6
N:	20169	18172	16079	20169	18172	16079
Third	Low	Middle	High	Low	Middle	High
Depvar:	$\Delta\ln(\text{wage})$	$\Delta\ln(\text{wage})$	$\Delta\ln(\text{wage})$	$\Delta\ln(\text{wage})$	$\Delta\ln(\text{wage})$	$\Delta\ln(\text{wage})$
t-1 $\Delta\ln(\text{offshore})$	0.003*** (0.000)	-0.007*** (0.001)	0.036*** (0.010)	-0.004 (0.003)	-0.001 (0.003)	-0.009** (0.004)
comp	0.062 (0.068)	-0.045 (0.083)	0.116*** (0.038)			
t-1 $\Delta\ln(\text{ofshr})$ *comp				0.048 (0.050)	-0.014 (0.067)	0.048 (0.093)
t-1 $\Delta\ln(\text{ofshr})$ *comm				0.152*** (0.050)	0.024* (0.008)	0.112* (0.063)
t-1 $\Delta\ln(\text{ofshr})$ *edu				0.002* (0.001)	0.001 (0.001)	0.001** (0.000)
t-1 $\Delta\ln(\text{ofshr})$ *nrtne				0.112 (0.096)	0.002 (0.010)	0.098** (0.048)
R-sq	0.21	0.19	0.22	0.26	0.15	0.26

*, ** and *** indicate significance at the 10, 5 and 1 percent levels

Robust standard errors are in parentheses

Table 7: Polarization within occupations

Model :	1	2
N:	65422	66359
Depvar:	$\Delta\ln(75/50)$	$\Delta\ln(50/25)$
t-1 $\Delta\ln(\text{offshore})$	0.001 (0.001)	-0.001 (0.001)
t-1 $\Delta\ln(\text{off})^*\text{comp}$	0.061** (0.026)	-0.069* (0.038)
t-1 $\Delta\ln(\text{off})^*\text{comm}$	-0.03 (0.027)	-0.021 (0.026)
t-1 $\Delta\ln(\text{off})^*\text{edu}$	0.0001*** (0.000)	0.0001*** (0.000)
t-1 $\Delta\ln(\text{off})^*\text{nrtne}$	0.028** (0.013)	0.069* (0.038)
R-sq	0.15	0.15

*,** and *** indicate significance at the 10, 5 and 1 percent levels. Robust standard errors are in parentheses

Table 8: Stage 1: Instrumental Variable Regressions

Model :	1
Level:	Industry
N:	532
Depvar:	$\Delta\ln(\text{offshore})$
EuroFDIChange	0.008** (0.003)
Δcomp	0.059 (0.040)
R-sq	0.02

*,** and *** indicate significance at the 10, 5 and 1 percent levels. Robust standard errors are in parentheses

Table 9: Stage 2: Industry Level Polarization

Model :	1	2	3	4	5	6
N:	530	530	517	517	504	504
Depvar:	$\Delta \ln(75/50)$	$\Delta \ln(50/25)$	$\Delta \ln(75/50)$	$\Delta \ln(50/25)$	$\Delta \ln(75/50)$	$\Delta \ln(50/25)$
Δcomp	0.005 (0.003)	0.052 (0.045)				
$\Delta \ln(\widehat{\text{offshore}})$	0.003 (0.005)	-0.014** (0.006)				
$t - 1\Delta \text{comp}$			0.005 (0.003)	-0.026 (0.029)		
$t - 1\Delta \ln(\widehat{\text{offshr}})$			0.022** (0.011)	-0.046*** (0.012)		
$t - 2\Delta \text{comp}$					0.007** (0.003)	-0.055** (0.026)
$t - 2\Delta \ln(\widehat{\text{offshr}})$					0.006** (0.003)	-0.001 (0.003)
R-sq	0.046	0.124	0.052	0.0403	0.111	0.115

*, ** and *** indicate significance at the 10, 5 and 1 percent levels

Robust standard errors are in parentheses

Table 10: Stage 2: Occupation Level

Model :	1	2
N:	45196	45196
Depvar:	$\Delta \ln(\text{wage})$	$\Delta \ln(\text{emp})$
$t - 1\Delta \ln(\widehat{\text{offshore}})$	-0.001* (0.001)	-0.032* (0.017)
$t - 1\Delta \ln(\widehat{\text{offshr}}) * \text{comp}$	-0.051** (0.026)	-0.646*** (0.137)
$t - 1\Delta \ln(\widehat{\text{offshr}}) * \text{comm}$	0.078*** (0.025)	0.985*** (0.329)
$t - 1\Delta \ln(\widehat{\text{offshr}}) * \text{edu}$	0.0005** (0.0002)	0.027*** (0.008)
$t - 1\Delta \ln(\widehat{\text{offshr}}) * \text{nrtne}$	0.075** (0.036)	0.038*** (0.011)
R-sq	0.12	0.13

*, ** and *** indicate significance at the 10, 5 and 1 percent levels

Robust standard errors are in parentheses

Table 11: Stage 2: Occupation Level IV by Thirds

Model :	1	2	3
N:	17246	17708	17914
Depvar:	$\Delta \ln(wage)$	$\Delta \ln(wage)$	$\Delta \ln(wage)$
Third	low	Middle	High
comp	0.135 (0.086)	-0.117 (0.076)	0.094** (0.047)
$t - 1 \widehat{\Delta \ln(offshore)}$	0.022 (0.021)	-0.110*** (0.021)	0.128*** (0.014)
R-sq	0.04	0.06	0.10

*,** and *** indicate significance at the 10, 5 and 1 percent levels
Robust standard errors are in parentheses

Table 12: Stage 2: Industry Level IV Polarization Controlling for Exports

Model :	1	2	3	4	5	6
N:	464	464	386	386	305	305
Depvar:	$\Delta \ln(75/50)$	$\Delta \ln(50/25)$	$\Delta \ln(75/50)$	$\Delta \ln(50/25)$	$\Delta \ln(75/50)$	$\Delta \ln(50/25)$
Δcomp	0.001 (0.001)	0.000 (0.001)				
$\Delta \ln(\widehat{\text{offshore}})$	0.026* (0.015)	-0.016** (0.007)				
$\Delta \ln(\text{export})$	0.002 (0.005)	0.001 (0.005)				
$t - 1 \Delta \text{comp}$			0.001 (0.002)	0.000 (0.000)		
$t - 1 \Delta \ln(\widehat{\text{offshr}})$			0.008* (0.004)	-0.045*** (0.014)		
$t - 1 \Delta \ln(\text{export})$			0.007 (0.007)	0.005 (0.005)		
$t - 2 \Delta \text{comp}$					0.001 (0.001)	-0.001* (0.001)
$t - 2 \Delta \ln(\widehat{\text{offshr}})$					0.004* (0.002)	-0.033** (0.015)
$t - 2 \Delta \ln(\text{export})$					0.004 (0.007)	-0.004 (0.007)
R-sq	0.02	0.01	0.01	0.05	0.02	0.02

*, ** and *** indicate significance at the 10, 5 and 1 percent levels

Robust standard errors are in parentheses

Table 13: Stage 2: Occupations Level IV by Thirds Controlling for Exports

Model :	1	2	3
N:	12823	9382	13323
Depvar:	$\Delta \ln(\text{wage})$	$\Delta \ln(\text{wage})$	$\Delta \ln(\text{wage})$
Third	Low	Middle	High
comp	0.235*** (0.046)	-0.079 (0.074)	0.165** (0.080)
$t - 1 \Delta \ln(\widehat{\text{offshore}})$	0.113 (0.116)	-0.023*** (0.008)	0.174*** (0.023)
$t - 1 \Delta \ln(\text{export})$	-0.001 (0.001)	0.001 (0.001)	0.002** (0.001)
R-sq	0.05	0.07	0.10

*, ** and *** indicate significance at the 10, 5 and 1 percent levels

Robust standard errors are in parentheses

Figure 1: Log hourly earnings ratios at the top and bottom of the wage distribution




Figure 2: Log hourly earnings change from 2001 to 2010 by percentile of the income distribution





Figure 3: Communication Task Intensity by 2002 Wage Quintile**Figure 4: Nonroutine Task Intensity by 2002 Wage Quintile**