The great reversal in the demand for skill and cognitive tasks

Paul Beaudry*
David A. Green†
Ben Sand‡

November 2013

Abstract

What explains the current low rate of employment in the US? While there has been substantial debate over this question in recent years, we believe that considerable added insight can be derived by focusing on changes in the labor market at the turn of the century. In particular, we argue that in about the year 2000, the demand for skill (or, more specifically, for cognitive tasks often associated with high educational skill) underwent a reversal. Many researchers have documented a strong, ongoing increase in the demand for skills in the decades leading up to 2000. In this paper, we document a decline in that demand in the years since 2000, even as the supply of high education workers continues to grow. We go on to show that, in response to this demand reversal, high-skilled workers have moved down the occupational ladder and have begun to perform jobs traditionally performed by lower-skilled workers. This de-skilling process, in turn, results in high-skilled workers pushing low-skilled workers even further down the occupational ladder and, to some degree, out of the labor force all together. In order to understand these patterns, we offer a simple extension to the standard skill biased technical change model that views cognitive tasks as a stock rather than a flow. We show how such a model can explain the reversal in the data that we present, and offers a novel interpretation of the current employment situation in the US.

*University of British Columbia, and NBER
†University of British Columbia and Research Fellow, IFS, London
‡York University.

Introduction

The poor performance of the US labor market since the financial crisis of 2008 continues to generate enormous debate fuelled by ongoing disagreement about the fundamental forces driving the decline. One side argues that difficulties in the US economy can be characterized as stemming from very weak aggregate demand due to a deleveraging process and the zero lower bound constraint on nominal interest rates. The other side argues for a set of explanations which emphasize more structural difficulties in combination with cyclical elements. Many of these structural explanations view the housing boom of the 2000s as having hidden (for a few years) longer term problems with employment options. The housing bust accompanying the cyclical downturn is then viewed as unveiling those problems. For example, the papers by Charles, Hurst, and Notowidigdo (2012) and Siu and Jaimovich (2012) emphasize the ongoing role of declining manufacturing employment and the disappearance of other routine jobs in causing the current low rates of employment. This type of explanation traces the origins of current difficulties back several decades, with a possible acceleration of the process more recently. While these different perspectives contain important elements for understanding the current situation, we believe that an important factor has been missed in the current debate.

In this paper, we argue that in about the year 2000, the demand for skill (or, more specifically, the demand for cognitive tasks that are often associated with high educational skill) underwent a reversal, and that this reversal can help in understanding poor labor market outcomes after 2000 more generally.[1] Numerous researchers have documented a substantial growth in the demand for occupations involving cognitive tasks (whether occurring exogenously (Katz and Murphy, 1992) or endogenously (Beaudry and Green, 2005; Acemoglu, 2002)) and an accompanying reduction in the demand for more middle-wage routine occupations in the two decades or so before 2000 (Juhn (1999); Autor, Levy, and Murnane (2003); Autor and Dorn (2013); Autor, Katz, and Kearney (2006, 2008); Dustmann, Ludsteck, and Schönberg (2009); Firpo, Fortin, and Lemieux (2011); Goos and Manning (2007)). This period is often described as a period of polarization, with an increased concentration of employment in either high-paying cognitive occupations or in lower-paying manual-service jobs. Since more educated workers have a comparative advantage in performing cognitive

[1] Throughout this paper, we will focus on three broad occupation groups that are based on the discussion in Acemoglu and Autor (2011). Cognitive task occupations consist mainly of managers, professionals and technical workers, and are seen as complementary to Information Technology capital and the organizational forms that go with it. Routine tasks are mainly production and clerical workers, and are seen as easily replaced by the new technology. Manual tasks are laborer and service occupations which require low skill but are not easily substituted for with IT capital.

tasks, this explains a large fraction of the increase in returns to education over the period. While the existing literature appears to reflect an implicit belief that this process has continued unabated in the period since the turn of the century, the first object of this paper will be to document that the demand for cognitive tasks has actually been declining since 2000. While such a decline in demand has had, and continues to have, a direct impact on more skilled workers, we go on to show that it has likely had a substantial impact on less skilled workers as well. In particular, we argue that in response to the demand reversal, high-skilled workers have moved down the occupational ladder and have begun to perform jobs traditionally performed by lower-skilled workers. This de-skilling process, in turn, resulted in high-skilled workers pushing low-skilled workers even further down the occupational ladder and, to some degree, out of the labor force all together. This process had been going on since 2000, but, as argued in earlier papers, the housing boom between 2003 and 2006 masked some of the effects which only become fully apparent after the financial crisis.

The contributions of this paper are twofold. First, we present a simple framework clarifying why skilled-biased technological change can cause a boom and bust in the demand for cognitive tasks. The key idea is that the IT revolution, and the revolution in organizational form that has gone with it, can be seen as a General Purpose Technology (GPT) and, like all GPTs before it, it will eventually reach maturity. If the implementation of the GPT has a capital investment form and cognitive tasks are a key component of the investment phase, under reasonable conditions, demand for cognitive tasks will have an over-shooting property. During the key investment stage, there will be high and growing demand for cognitive tasks to build the new capital, but once the new capital is in place, cognitive task workers are only needed to maintain the new capital. At this maturity stage, there will be greater demand for cognitive tasks than before the technological revolution but we will see a reduction in demand for these tasks relative to the peak investment stage. We argue that the turn of the century is that approximate turning point from the peak investment to the maturity stage. Importantly, it is not the case that all innovation related to the GPT needs to cease in the maturity stage for this model to fit the data patterns we describe; only that it slow down. As an historical example, innovation related to electricity continued long after the investments involved in building the spine of the national electrical system were completed.

In describing the adjustment process for this cycle, we exploit insights of the extensive task-versus-skill literature (see Acemoglu and Autor (2011) for an overview) which emphasizes that changes in task demands will affect workers across the entire labor market – not just those currently performing these tasks – since workers will adapt to changing circumstances by redirecting the supply of their skills to the other tasks they can perform. In particular, we are arguing that relative to the 1990s, the post 2000 maturity era for the IT revolution is one where even the demand for skilled workers is reduced. In this maturity stage, having a college degree is only partly about obtaining access to high-paying managerial and technology jobs since it is also about beating out less educated workers for barista and clerical job type jobs.

Our second main contribution is that we provide a detailed picture of the changes in employment patterns and wages over the last thirty years with a particular focus on the reversal in the growth in demand for more cognitive intensive occupations and the adjustment of more skilled workers to this change. Importantly, we view this contribution as standing even if the reader is not convinced about our specific story behind the cognitive task demand reversal.

The remaining sections of the paper are structured as follows. In section 1 we begin by presenting some very broad labor market trends which highlight the salience of the year 2000 as an important turning point in the US economy. Since 2000 is the year of the Tech Bust, this motivates some of our modelling choices in Section 2. In that section, we present a simple dynamic model of adjustment to new technological opportunities which creates a cycle in the demand for cognitive tasks together with a continuous decline in the demand for routine tasks. During this process, workers with different skills shift their supply of labor across tasks as a means of adapting to the changes in demand. Since the model is highly stylized, at the end of the section we provide an heuristic generalization which provides an general framework for exploring the empirical relevance of our story. Section 3 looks at a large set of data patterns. In particular, we examine the employment patterns in different sets of occupations, the assignment of workers of different skills to tasks, and adjustment in wages. A key challenge in our empirical investigations will be to try to focus on skill price changes by netting out changes in the composition of the labor force arising, for example, from increased educational attainment in the population as a whole. We focus much of our attention on wage adjustments of younger workers as we believe these best reflect current changes in market forces.

1 Aggregate Employment and Average Real Wages

1.1 Data

We begin our investigation by presenting some key labor market trends for the US. The data we use for this (and the empirical exercises later in the paper) are drawn from the Outgoing Rotation Group (ORG) Current Population Survey Supplements for the years 1979-2011. Following Lemieux (2006), we use the hourly wage as our wage measure, weight observations by hours worked in combination with the ORG weights, and do not use observations with allocated wages.[2] Wages, hours of work and employment status refer to the week prior to the survey week. We present annual values by averaging across all months in a calendar year. For our employment rate constructs, we sum the number of respondents who report working in the reference week over the calendar year and divide this by the sum of working age respondents in the calender year. In doing so, we adjust the ORG weights such that the annual sum equals the size of the US population for a particular group. In all our empirical work, we restrict the sample to individuals aged 18-64 with positive potential work experience.

1.2 Wage and Employment Rate Patterns

In Figure 1 we plot the employment rate of individuals 18-64. We superimpose on the figure an estimated linear trend allowing for one break. The pattern in the figure is quite clear and rather well known. The US employment rate increases relatively continuously over the 1980s and 1990s, and then this growth reverses around 2000 (testing for the single optimal break point actually indicates 1999 as the point of the break with these data). The growth in the employment rate in the 1980s and part of the 1990s is dominated by the trend increase in participation of women in the labor force. It is striking to note that the reversal after 2000 was so strong that by 2010 the employment rate was back to a level close to what was recorded in 1980. In Figure 2 we plot the average real hourly wage for all workers against the employment rate. We focus on the post 1990 period to focus attention away (to some extent) from the role of the increased labor attachment of women. Figures 3 and 4 contain the same wage-employment diagram for men and women separately. All three figures paint

[2] Toward the end of the paper, when we examine percentiles of the wage distribution we use both allocated and non-allocated wages. The Data appendix provides additional information on our data processing.

a similar picture: from 1990 to around 2000, both employment rates and average real wages rose. Then after 2000, employment rates started to decline while average wages did not. Moreover, in each of these figures, the 2003-2007 period appears as what might be characterized as a stalling phase in a more general process.

One might suspect that the increase in average wages observed after 2000 in Figure 2 simply reflects a shift in the composition of employment toward more highly educated workers. However, as we show in the appendix, when we plot the analogue of that figure for high school and college educated workers separately, we see a very similar pattern for each group. So, if the behavior of average wages after 2000 is driven by some sort of selection process, it is most likely within-education groups. Nonetheless, there is one important cross-education group distinction which we highlight in Figure 5. In this figure, we plot the ratio of the employment rate of high school educated workers to that of college educated workers. From this we see that, prior to 2000, the employment ratio of the high school workers is increasing faster than that of the college workers, while after 2000 we have the reverse. Thus, changes in the employment rate of less educated workers play an important role in the employment rate patterns observed in Figure 2.

The patterns in these first five figures motivate us to downplay the housing boom and bust of 2003-2008 for understanding the current labor market situation and instead incites us to focus on more long-run trends. In particular, we want to ask: "What type of force could be driving such a medium run process where employment and wages increase for at least a decade follwed by a reversal?" Since 2000 was the year of the bust in the high-tech sector, it seems reasonable to us to begin by exploring this issue with a model which emphasizes technological change as presented in the following section. In the choice of modelling strategy, we are aiming at a unified explanation of both the pre-2000 and post-2000 period in order to offer, among others, a qualitative and unified explanation to the inverse 'C' relationship between employment rates and average wages presented in Figure 2. The model we present will have a set of implications for employment, wage behavior and occupational choice; and we will show how the data support these implications.[3]

To further motivate our focus on explaining current labor market patterns as reflecting a medium run boom and bust driven by the IT revolution, in Figures 6

[3] We should note from the outset that the model we present is not the only one capable of rationalizing the data patterns we emphasize. We see several alternative avenues for explaining the patterns presented in this paper. For example, we could adopt an approach with two different driving forces, one that explains the pre-2000 period and one that explains the post-2000 period, and the data we present may not be able to differentiate the two views. Nonetheless, we think the model is very useful in helping organize an approach to the data.

to 9 we report a set of changes in the share of investment in information technology over the last few decades. Figure 6 reports the share of GDP directed to investment in equipment and software. As can be seen, this share rose by close to 40% from the mid-70s to 2000, with most of that rise occurring in the 1990s. However, since 2000 this share has been on a downward trajectory, bring us back to levels of the mid-seventies. Figures 7, 8 and 9 echo the same message using more restrictive measures of investment in IT. Figure 6 plots investment in information processing equipment and software as a ratio to GDP, while Figure 8 focuses only on investment in computer hardware. In both these figures we see a more than 50% rise in the investment share over the 80s and 90s, followed by a drastic and sustained fall post 2000. Figure 9 reports the investment share in software. Here the pattern is less dramatic. Software investment also experienced a huge increase in the 80s and 90s, but since 2000 there has essentially seen a stagnation for this series instead of a fall. All these figures support the notion that there appears to have been an important change around the year 2000. Our main claim is that the same force driving the investment patterns can help in understanding labor market patterns.

2 A model of boom and bust in the demand for cognitive tasks

The goal of this section is to show how a rather standard model of skill-biased technological change, extended to include a dynamic adjustment process, can create a boom-bust cycle in the demand for cognitive tasks along with a continuous decline in the demand for routine tasks. Our aim will be to emphasize simple empirical implications of the model, which can then be readily compared with data.

2.1 Basic Model

Consider an environment with three types of agents: highly educated individuals, less educated individuals, and entrepreneurs who run firms. There is one consumption good which plays the role of the numeraire. All individuals are risk neutral and discount the future at rate, ρ. The entrepreneurs hire individuals to perform two distinct tasks. One task will be referred to as a cognitive task (or cognitive occupation) and one as a routine task. Individuals can perform only one task at a time, choosing where to supply their labor based on comparative advantage. The market for each task is assumed to be competitive and to function in a Walrasian fashion. The production possibilities available to the entrepreneur will vary over time to reflect technological change in favor of the more cognitive task. Our main departure from the literature on skill-biased technological change will be the way we assume the cognitive task affects production. In particular, instead of assuming that these tasks only affect current production, we view these tasks as building intangible capital for the firm in the form of organizational capital denoted by Ω_t. We refer to this as organizational capital in order to emphasize recent changes that go beyond the direct use of computers in production. We will show that this simple alteration leads to a model with a boom and bust in demand for the cognitive task.

Defining L_t^c as the effective units of the cognitive task hired by the representative firm and L_t^r as the effective units of the routine task, we can represent the optimization problem faced by the entrepreneur as choosing L_t^c and L_t^r to maximize profits given by,

$$\max_{\{L_t^c\},\{L_t^r\}} \int_0^{\infty} [F(\Omega_t, L_t^r, N, \theta_t) - w_t^c L_t^c - w_t^r L_t^r] exp^{\rho t} dt$$

s.t

$$\dot{\Omega} = L_t^c - \delta\Omega, \tag{1}$$

where $F(\cdot)$ is the instantaneous production function of the consumption good, N is the entrepreneur's time endowment, θ_t is an technology parameter, δ is the depreciation rate of organizational capital, w_t^c is the price of an effective unit of cognitive skill and w_t^r is the price of an effective unit of routine skill. The production function is assumed to be increasing in all its arguments, concave, and to exhibit constant returns to scale with respect to L^c, L^r, and N. For simplicity, we will normalize $N = 1$ and drop it from further notation. The law of motion for organizational capital, Ω, which says that it is created with cognitive tasks and depreciates at a constant rate, will play a key role in the insights gained from the model.

We make two main assumptions regarding the production function. The first (which is standard in the large literature on skill-biased technical change and polarization) is that the organizational capital produced by cognitive tasks is a substitute for routine labor, that is, $F_{\Omega,L^r} < 0$. Second, we assume that an increase in θ increases the productivity of the organizational capital produced by the cognitive task, that is, $F_{\Omega,\theta} > 0$. For simplicity, we also assume that θ has no direct effect on the productivity of routine tasks, that is, $F_{L^r,\theta} = 0$.[4] Many examples of this type of technology (including the use of computers in clerical workplaces) have been presented in discussions of technological change and polarization.

[4]We could relax this last assumption and allow for a more general production technology. What we need is that the main effect of technological change is to raise the productivity of Ω.

The first order conditions associated with this optimization are given by:

$$F_{L^r}(\Omega_t, L_t^r, N, \theta_t) = w_t^r \tag{2}$$

$$\dot{w}_t^c = (\delta + \rho)w_t^c - F_\Omega(\Omega_t, L_t^r, N, \theta_t). \tag{3}$$

The first condition indicates that units of the routine task should be hired up until their marginal product is equal to their task price, while the second condition indicates that organizational capital should be accumulated to the point where its marginal product is equal to its user cost inclusive of capital gains of losses on the value of Ω (note that the shadow price of Ω is simply w_t^c). The two conditions, (2) and (3), combined with the accumulation equation, $\dot{Q} = L_t^c - \delta Q$, implicitly define the demands for cognitive and routine tasks as functions of their prices.

To complete the model we need to determine how the supplies of cognitive and routine tasks respond to prices. This requires specifying the labor supply decisions of the the high- and low-educated workers. The problem faced by these individuals will be virtually identical except for the fact that the distribution of their relative productivities in the two tasks will differ. Let us begin by considering the decision problem faced by a high-education worker.

Assume there is a measure H of highly educated workers and each of these individuals has an identifier ψ drawn from a uniform distribution defined over the unit interval $[0, 1]$. The identifier ψ gives their productivity rank in cognitive tasks with a function, $h(\psi)$, translating that rank into effective units of cognitive skill. For convenience, we will define $h(\psi)$ such that $h'(\psi) \leq 0$, i.e., ψ ranks individuals in decreasing order of productivity. Accordingly, if a high educated individual indexed by ψ decides to supply her labor to the cognitive occupation, she will receive a wage, $h(\psi)w_t^c$. In contrast, we will assume that her productivity in routine tasks is independent of ψ and is such that she would supply $1 + \alpha$ effective units of the routine task, (where $\alpha > 0$). This implies that her wage payment if she supplied her labor to the routine task would be $(1 + \alpha) \cdot w_t^r$. The individual also has the option of home production, where her labor can produce A units of the consumption good. As will become clear, in equilibrium w_t^r will be greater or equal to A so that highly educated workers will not choose to stay in the home sector. Thus, their optimal decision will be characterized by a cut-off for ψ denoted $\bar{\psi}_t$ and defined by:

$$w_t^c \cdot h(\bar{\psi}_t) = (1 + \alpha) \cdot w_t^r, \tag{4}$$

with individuals having a $\psi \leq \bar{\psi}_t$ supplying their labor to cognitive tasks and those with $\psi > \bar{\psi}_t$ supplying their labor to the market for routine work.[5]

The decision problem for less educated workers is very similar. There is a measure W of the less educated workers, and each of these has an index ϕ drawn from a uniform distribution on $[0, 1]$, with $g(\phi)$ giving the effective units of cognitive skill of a less educated worker with rank ϕ. The first difference between the more and less educated workers is that $h(x) > g(x)$ for all $x \in [0, 1]$; that is, less educated workers of each rank generate fewer effective units of cognitive tasks. For simplicity, we will assume that these effective unit functions have the same shape, with $g(x) = b \cdot h(x)$ and $0 \leq b < 1$. If a less educated worker supplies his labor to routine tasks, he will receive a wage payment w_t^r since we assume that his labor is equivalent to 1 effective unit of routine work independent of his rank, ϕ. The fact that the effective labor of less educated workers is also lower in routine tasks is the second difference that separates the two classes of workers.[6] As with the more educated workers, less educated workers produce A units of the consumption good if they choose home production.

The decision problem of the less educated workers is characterized by a cut-off $\bar{\phi}_t$, implicitly defined by:

$$w_t^c \cdot b \cdot h(\bar{\phi}_t) = w_t^r, \tag{5}$$

where all workers with $\phi \leq \bar{\phi}_t$ supply their labor to cognitive tasks. For workers with $\phi > \bar{\phi}_t$ there are two possible equilibrium configurations. Either $w_t^r > A$ and they supply their labor to routine tasks, or $w_t^r = A$ and they are indifferent between working at routine tasks or staying at home. In the latter case, the division of labor between routine jobs and home work will be determined solely by demand. In particular, if $w_t^r = A$, the employment of routine tasks is implicitly determined by $F_{L^r}(\Omega_t, L_t^r, N, \theta_t) = A$. We will assume that routine employment for less educated workers is this demand minus the amount of the routine task supplied by more educated workers.[7] As a result, the fraction of less educated workers in routine jobs at a point in time will be given by:

$$\frac{L_t^r - H(1 - \bar{\psi})(1 + \alpha)}{W}.$$

Using the decision rules for both classes of workers, the total supply of effective

[5]If $\bar{\psi}_t \geq 1$, then all high-education workers supply their labor to cognitive occupations.

[6]This could occur, for example, even if more educated workers are not more productive per unit time working at a routine task but are more likely to show up to work on time each day.

[7]In order to guarantee that more educated worker have a comparative advantage in cognitive jobs, we assume that $b \cdot (1 + \alpha) < 1$.

Table 1: Summary of Key Model Constructs

	Cognitive Emp.	Cognitive Wage	Routine Emp.	Routine Wage
High Education	$H \cdot \bar{\psi}_t$	$w_t^c \int_0^{\psi_t} \frac{h(\psi)}{\bar{\psi}_t} d\psi$	$H(1-\bar{\psi}_t)$	$(1+\alpha)A$
Low Education	$W \cdot \bar{\phi}_t$	$w_t^c \int_0^{\phi_t} \frac{h(\phi)}{\bar{\phi}_t} d\phi$	$L_t^r - H(1+\alpha)(1-\bar{\psi}_t)$	A
Emp. Rate	$\frac{H \cdot \bar{\psi}_t + W \cdot \bar{\phi}_t}{H+W}$		$\frac{L_t^r - \alpha H}{H+W}$	

units of cognitive task labor can be expressed as:

$$L_t^c = H \int_0^{\bar{\psi}} h(\psi) d\psi + W \int_0^{\bar{\phi}} g(\phi) d\phi, \tag{6}$$

where the upper limits of integration are determined by the equations (4) and (5).[8]

As we want to focus on a case where technological adjustment can change the economy's employment rate, we will assume that parameter values are such that equilibrium is characterized by $w_t^r = A$ and the employment rate becomes demand determined. In this case, the equilibrium determination of $L_t^c, L_t^r, \Omega_t, w_t^c, \bar{\psi}_t, \bar{\phi}_t$ is given by the solution to the system of equations (1) through (6). Once we know these quantities, we can easily derive – among others – the number of workers in each occupational-education cell and the average wage for each of these cells, as well as the employment rate for different occupations or different education groups. Table 1 provides explicit expressions for some key employment and wage entities.

2.2 Steady State Implications

We now turn to the question of how such an economy reacts to an increase in θ. In particular, we want to highlight the dynamic properties of this model when, starting from a steady state with constant θ, there is an improvement of θ over time which takes the shape of a diffusion process. However, before looking at the dynamic properties, we want to highlight the difference between an initial steady state with $\theta = \theta_0$ and a later steady state with $\theta = \theta_1 > \theta_0$. In doing this, we want to emphasize that when looking only at steady state comparisons, the model maintains the key features emphasized in the skill biased technological change literature. This is stated in Proposition 1.

Proposition 1. The steady state effects of an increase in θ are:

[8] To be more precise, the upper limits of integration should be of the form $\max[1, \bar{\psi}]$ and $\max[1, \bar{\phi}]$.

– An increase in the employment rate in cognitive occupations and a decrease in the employment rate in routine occupations.

– Skill upgrading in the sense that, for each education class, a greater fraction of individuals is in cognitive occupations.

– The wage differential between more and less educated workers will increase as long as b is not too big.[9]

Proposition 1 indicates that when comparing steady states, our model mimics implications of existing models of skill biased technological change with endogenous occupational choice (See, for example, Acemoglu and Autor (2011)). However, the addition of organizational capital to the standard model does permit some extra steady state implications, as summarized in the following proposition:

Proposition 1B If δ is sufficiently small, the steady state effects of an increase in θ are:

– The overall average wage increases,

– The overall employment rate decreases.

The intuition behind this proposition is that when δ is small, the fruits of cognitive employment act like a capital stock. With more organizational capital being used in the higher θ steady state, more cognitive tasks are needed to offset depreciation. This, in turn, implies more workers in the high wage, cognitive occupation and, thus, a higher overall average wage. However, with more organizational capital in the new regime, there is also lower demand for routine tasks and if the number of added workers needed to service the new organizational capital is not too large (i.e., if δ is small) then the net effect is a reduction in demand for labor overall, and a reduction in the employment rate. This fits with the long run pattern shown in our first set of figures, with a decrease in the employment rate and an increase in the average wage between the early 1990s and the late 2000s. While at first pass, such a pattern may seem to require a inward shift of a labor supply curve, our model suggests it could be the result of technological change. To explore the relevance of such a possibility, we now examine the model's dynamic implications.

[9] In this model it is possible for an increase in θ to reduce the differential between more and less educated workers. This arises when the technological change pushes a sufficiently large number of the lowest paid less educated individuals out of market employment, leaving mainly the lower educated workers employed in relatively high paying cognitive jobs. However, this will not arise if b is sufficiently small.

2.3 Dynamics

The new insights of our model come from its dynamic implications for an economy adjusting to the diffusion of technological knowledge. For simplicity, we take an extreme form of a diffusion process whereby at time 0 the economy is in a steady state with $\theta = \theta_0$, and then all the agents in the economy learn that θ will increase to $\theta_1 > \theta_0$ at time τ. Thus, the diffusion process used here is a step function, which allows for an easy characterization of the problem. The results extend easily to a more gradual process, including one where diffusion is first rapid then slows down as the technology gradually reaches maturity.

To analyze this dynamic system, we can either examine the issue numerically, take a linear approximation of the system, or adopt simple functional forms. We choose to follow the latter route as it allows for an analytical solution and simplifies the exposition. Thus, we will assume that the production function takes the following quadratic form: $F(\Omega_t, L_t^r, \theta_t) = \gamma_1\theta_t\Omega_t + \gamma_2 L_t^r - \gamma_3\Omega_t L_t^r - \gamma_4\Omega_t^2 - \gamma_5(L_t^r)^2$, with the concavity restrictions $(4\gamma_4\gamma_5 > \gamma_3^2)$. For the functions determining the supply of effective units of cognitive labor, we assume that $h(\psi) = \psi^{-\frac{1}{2}}$ and $g(\phi) = b\phi^{-\frac{1}{2}}$, with $0 \leq b < 1$. Under this parametrization, our dynamic system can be reduced to a two dimension linear system in L_t^c and Ω_t as given by (1) and:

$$\dot{L}^c = (\delta + \rho)L_t^C + \frac{4\gamma_4\gamma_5 - \gamma_3^2}{2\gamma_5}\Omega_t - \frac{\gamma_1}{\eta}\theta_t + \kappa, \tag{7}$$

where $\eta = \frac{2(H+W(1+\alpha)b^2)}{(1+\alpha)A}$ and $\kappa = \frac{\gamma_3(\gamma_2 - A)}{2\gamma_5}$.[10] Given a solution for this system, the values of $L_t^r, w_t^c, \bar{\phi}_t$ and $\bar{\psi}_t$ are then given by:

$$L_t^r = \frac{\gamma_2 - \gamma_3\Omega_t - A}{2\gamma_5},$$

$$W_t^c = \frac{1}{\eta}L_t^c,$$

$$\bar{\psi}_t = \left(\frac{w_t^c}{(1+\alpha)A}\right)^2,$$

$$\bar{\phi}_t = \left(\frac{bw_t^c}{(1+\alpha)A}\right)^2.$$

The dynamics of this two variable system can be represented by a phase diagram as given in Figure 10. In this figure we see that there is an initial stage where the

[10] The two boundary conditions for this system are an initial value for Ω_t and the transversality condition associated with the entrepreneurs' optimization problem. Throughout, we assume that $\gamma_2 > \gamma_3\Omega_t + A$ so as to have positive employment in routine jobs.

employment of cognitive skills increase followed by a period of decrease. Throughout this process, Ω_t increases and the employment of routine skills decrease. The turning point for cognitive skill employment is at time $t = \tau$ where the technology becomes fully operational.[11] We present the implications for the economy in two propositions, with Proposition 2 containing the implications during the period when cognitive task employment is increasing and Proposition 3 corresponding to the period when it is decreasing.[12]

Proposition 2. Upon learning of the diffusion process for technology, the economy will initially go through a stage where:

– The average employment rate and, by implication, the employment rate of less educated workers, will increase.

– The average wage will increase, as will the average wage of each educational class.

– There will be skill upgrading in the sense that the fraction of employment in cognitive jobs of both education groups will increase.

Proposition 2 indicates that the arrival of the new technological opportunities will lead to an initial period where the economy can be unambiguously described as a going through a boom. In particular, the economy will initially experience increased employment and wages, and this is beneficial to both types of workers. Intuitively, during the boom period there is increasing demand for cognitive tasks in order to build the organizational capital that will allow the economy to take full advantage of the technological change. This will generate increases in the cognitive task price, w_t^c, which will draw both high- and low-education workers into cognitive occupations. This, in turn, puts pressure on the routine task market that draws more individuals out of the home production sector, raising the employment rate.

However, as the next proposition indicates, this boom will eventually by followed by a bust period. It is this subsequent bust period induced by the process of technology adoption which is the key insight of the model.[13]

[11] Note that these dynamics mimic those associated with a standard Q-theory of investment with an anticipated shock.

[12] Here we do not need to assume that b and α are small.

[13] There are several models in the literature which suggest that skill biased technological change can create an initial bust period (e.g., Carlaw and Lipsey (2002)), but seldom do they predict a later bust period.

Proposition 3. Due to the diffusion of the new technology, the economy will eventually go into a bust phase (which will last until the new steady state is attained) with the following properties (expressed relative to the boom period):

– A decrease in the aggregate employment rate including a decrease in the employment rate of both cognitive and routine occupations.

– Skill degrading in the sense that, for each education class, a lower fraction of individuals is employed in cognitive occupations.

– For less educated workers, there is also a reduction in the fraction of individuals in routine tasks and an increase in non-employment.

– Except during a bust, there are continued increases in returns for the entrepreneurial class.

Although technological change in the model has only positive impacts on the productivity of cognitive tasks, Proposition 3 states that this change eventually leads to a period characterized by a decreasing path for the cognitive task employment rate. This arises because cognitive tasks are modelled as creating a stock of organizational capital for the firm. Hence, there is an initial period when firms want to hire cognitive task workers to increase the stock of this capital, but eventually, when the stock is sufficiently large, there is less need for cognitive employment as it is only required to offset depreciation. If δ happens to be very close to zero, then the change in employment in cognitive tasks induced by a change in θ would be essentially an entirely temporary phenomena. Once the economy enters into the period of declining cognitive task employment, it is clearly in a bust period as it is also the case that employment in routine tasks continue to be replaced by the improved organizational capital.

It is interesting to highlight the skill downgrading process and its crowding out effect of the employment of low-educated workers during the bust period. The reduction in demand for cognitive jobs during this period implies that high-educated workers switch, in part, to accepting routine jobs. This movement of high-educated workers into the less skilled occupations amplifies the push of less educated workers toward non-employment. In fact, less educated workers move out of cognitive jobs because of the decrease in demand for those tasks, and they move out of routine jobs both because of decreased demand and because of increased supply to those jobs by the higher educated individuals. In this sense, employment has what we think of as a cascading nature, with more skilled workers flowing down the occupation structure and pushing less skilled workers even further down. Hence, even though the major change in the bust period relative to the boom period is the shift in the demand for cognitive jobs, non-employment increases among the less-educated as this is the main escape valve for the labor market.

Because of some of the stark assumptions we use in the model, its implications for wages are more stylized than those for quantities. In particular, it can easily be verified that the average wage for each education-occupation grouping is actually a constant in this model.[14] For the routine occupations this is not surprising since wages are pinned down by the value of home production, which is a constant by assumption. For the cognitive occupation this may be somewhat more surprising since the price of effective units of cognitive skill goes though an identical boom-bust cycle to that of employment in cognitive occupations. However, selection acts to decouple movements in the skill price from movements in the observed average wage in cognitive occupations. When the price of effective units goes up, the marginally less productive individuals enter these occupations, bringing down the average observed wage. Given the functional forms assumed for task-generation functions, these two effects cancel each other out, leaving average wages within occupation-education classes constant. Thus, the model's insights for occupational average wages is limited, and one needs to focus either on trying to obtain estimates of skill price movements by controlling for selection or on average wages for skill groups. For groups defined by skill, selection is assumed not to be an issue and thus wage implications can be examined directly.

With these selection forces in mind, it is interesting to consider how the average wage for the whole economy behaves during the bust period.[15] On the one hand, during this period there is a movement away from the high-paying cognitive jobs to the low-paying routine jobs, and this should depress the average wage. However, at the same time, less educated workers are leaving the labor force to non-employment, and since the departing individuals had the lowest wages, this tends to increase the overall average wage. Hence, the bust in this model – although it is spread widely across the economy – can be consistent with an average wage that looks rather unresponsive to the decline in employment.

To illustrate this latter fact, Figure 11 plots the joint movement of the average wage and the employment rate for a simple parametrization of the model.[16] As can

[14] For example, the average wage for high educated workers in the cognitive occupation is $2A(1+\alpha)$, while for the lower educated individuals in the same occupation it is simply $2A$.

[15] Proposition 1 already indicated that it will initially increase during the boom period.

[16] Figure is generated from a discrete-time version of the model. The production function is given by $F(\Omega, L^r, \theta) = \alpha_1 Q + \alpha_2 L^r - \alpha_3 Q L^r - \alpha_4 Q^2 - \alpha_5 (L^r)^2 + \alpha_6 \theta Q$, with parametrization $(\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6) = (2, 4.3, 1.1, 0.4, 1.2, 1)$. The firm's discount factor is $\beta = 0.95$, and the depreciation rate on organizational capital is $\delta = 0.1$. A highly educated individual of rank $\psi \in [0, 1]$

be seen in the figure, the employment rate and the average wage initially increase and then there is a reversal with the employment decreasing but the average wage actually increasing slightly. While this figure should be taken only as illustrative since the parameters are chosen arbitrarily, it does convey that this model can qualitatively replicate the pattern reported in Figure 2. More specifically, when Propositions 2 and 3 are combined with Proposition 1B we see that the model offers an explanation for the patterns described in Figures 2-4. Recall that in the 1990s, these figures depict an economy with an increasing average wage and rising employment rates - exactly what is predicted for the economy in Proposition 2 if we see the 1990s as part of the boom period. In contrast, the post-2000 period fits with the Proposition 3 prediction of declining employment rates and with the claim that selection implies muted movements in the overall average wage. Proposition 1B states that if δ is sufficiently small then the process will converge to a point where the employment rate is lower and the average wage higher than the point of departure, which fits with the long term differences in these figures. Thus, the model provides a unified explanation for the wage and employment patterns in the 1990s and 2000s with only one driving force. In contrast, an explanation built from a simple demand and supply model would require both positive movements in demand in the 1990s and negative supply shifts in the 2000s. Our model avoids the need to invoke unexplained labor supply disturbances to understand the patterns.

2.4 Main Implications and Manual Tasks

The model presented to this point is useful in terms of providing a way to organize our thoughts about empirical patterns (e.g., highlighting difficulties in interpreting average occupational wages) but it is also highly stylized. In this subsection, we recast the key elements and outcomes of the model in a simple demand and supply framework. Our goal is to try to emphasize the key lessons and direct attention away from simplifying assumptions that were made for reasons of tractability and logical completeness.

One key way in which the model is stylized is in its use of only three sectors

working in the cognitive sector supplies $h(\psi) = a\psi^{-1/2}$ efficiency units of labor, while a less-educated individual of rank $\phi \in [0, 1]$ supplies $g(\phi) = b\phi^{-1/2}$ efficiency units of labor, where we set $a = 0.15$ and $b = 0.1$. In the routine sector, highly educated and less-educated workers supply $1 + \alpha$ and 1 effective units of labor, respectively, where $\alpha = 0.05$. All workers produce $A = 1$ units of output per unit of labor in the home sector. There are total measures $\mu_H = 0.7$ and $\mu_L = 1$ of highly educated and less-educated workers, respectively. The Figure plots the response of the economy to shocks to θ as follows: at $t = 0$, the economy is in steady state with $\theta_0 = 0.2$. At $t = 1$, agents learn that at $t = 10$, θ will increase to $\theta_1 = 0.3$ and remain at that level thereafter.

(cognitive, routine, and home). This omits the non-routine manual sector which has been the focus of a substantial portion of the literature on polarization. We made this choice because we do not have anything to add to what the existing literature has found for these jobs in the 1990s and including it in the model would not have changed the conclusions about the cognitive and routine sectors but would have made the model too cumbersome to provide clean insights into what we see as a new point: the demand reversal for cognitive tasks. In the more heuristic description in this section, however, we will add the manual task sector.

Figure 12a provides a depiction of what we see as the dominant event in the cognitive task market in the boom period (the 1990s): a shift out in demand.[17] In our model, this happens because firms are building organizational capital to take advantage of the technical change but it could also arise from a more standard model with cognitive tasks entering the production function directly and a technical change biased in favor of these tasks. Figure 12b shows the market for routine tasks in the 1990s. Here we see both a shift down in demand (since the technical change is biased against such tasks) and a shift back in supply as workers are drawn up into the cognitive task market. Again, both of these would be found in standard polarization models, though the supply shift tends not to be emphasized. Taken together, we expect to see a rise in employment and the skill price for cognitive tasks and a decline in employment and possibly the skill price for routine tasks.[18]

In Figure 13a, we represent the market for manual tasks in the boom period. We follow Autor and Dorn (2013) in viewing manual tasks as directly providing a service good to households. In the boom phase, the demand for such services is likely to increase because of increased demand for services coming from segments of the population that are becoming richer. This is captured in the figure by a shift up in the demand curve for manual tasks. Given that demand is increasing for cognitive tasks and decreasing for routine tasks during the boom phase, the net effect on the supply of workers to manual occupations is unclear. For this reason, we have kept the supply curve for manual tasks in Figure 13a unchanged. Accordingly, with an increased demand for manual work, we should expect the price of the task to rise, as shown in the figure. The increased price would then favor movement of people from the household sector to the manual task sector. This is depicted by the reduction in

[17]In this section, as in the model, we abstract from general shifts in the supply of high educated workers.

[18]Note that in both markets, on the x-axis we have the number of efficiency units and on the y-axis is the price of the task. As we emphasized earlier, it is important to bear in mind that the price of the task is not generally equivalent to the observed wage in a sector since wage payments are a combination of the price of the task and the number of efficiency units of the person performing the task.

the supply in the household sector in Figure 13b. Thus, this simple extension of the model captures a polarization of jobs during the boom phase with employment and the task price increasing in both the cognitive and manual sectors.

Figure 14a captures what we view as the main new insight in our paper: the idea that the demand for cognitive tasks shifted down after 2000 (the bust period). In our model, this happens because of the nature of adjustment of the organizational capital stock but, as we discussed earlier, it could also happen for other reasons such as the arrival of a new technology that is biased against (most) cognitive tasks. What we want to emphasize is less the reason for the reversal in our model than the fact of a reversal itself and its implications for the remainder of the labor market. Those implications are captured, in part, in Figure 14b, which represents events in the routine sector in the bust phase. In this sector, the technological change continues to shift the demand curve downward while the supply curve shifts out because the falling skill price in the cognitive sector induces workers (who are predominantly high educated) to shift toward other sectors. The result is a clear implication in terms of a falling routine task price and falling employment if the demand shift is stronger than the supply shift.

In Figure 15, we graph supply and demand in the manual and home sectors during the bust phase. With skill prices falling in both the cognitive and routine sectors, more skilled workers move down the occupation ladder. This de-skilling process will tend to increase the supply of workers to the manual task sector thereby putting downward pressure on the task price in that sector. With a depressed price now in all three task markets, this will also tend to push the least skilled individuals out of the market altogether and into the home sector, as represented by the shift out in the supply for labor in the home sector. The net results in the bust phase for this extended version of the model is decreased employment in both the cognitive sectors and the routine sector, increased employment in the manual sector, and decreased employment rates as the least skilled leave the market. In the cognitive task sector, we would expect to see a falling skill price along with declining employment, as fits with a decline in demand. In comparison, in the manual sector, we expect to see movements that fit with an outward shift in demand: more employment but lower skill prices. This contrasts with the evidence of a positive shift in demand in this sector in the 1990s documented by Autor and co-authors. Finally, in the routine sector, the combination of supply and demand shifts make implications somewhat less clear, but we expect to see a decline in the skill price and a decline in employment. Whether driven by our specific mechanism or not, the key points are: a) that a period of increased demand for cognitive tasks was followed by a period of declining demand; and b) that this has generated first a drawing of workers up the occupational skill ladder (and into the labor market in general) and then a cascading of workers down the skill ladder (and out of the labor market for some); and c) that together these movements can explain the patterns in our initial figures as being driven ultimately by the reversal in the cognitive task market. We investigate these broad patterns empirically in the next section.

Finally, it is interesting to consider the relative sizes of the declines in the skill prices in the three sectors during the bust phase. It is not uncommon to depict the routine sector as an imperfectly competitive sector (perhaps because of unions) with workers being paid above market clearing wages. In that case, the workers entering the manual job sector in the bust phase would not be indifferent between manual and routine jobs, and we would expect to see relatively larger declines in manual occupation wages.

3 Patterns of Employment, Skill Upgrading and Task Prices

In this section we use data from the US Current Population Survey (CPS) from 1980-2010 to document three sets of labor market patterns which we argue support the simple boom-bust model of technological change we outlined in Section 2. We begin by examining aggregate changes in employment across occupations, focusing on differences between the 1990s and the 2000s. Then we turn to the changes in the distribution of workers across occupations, conditioning on worker skill. Finally, we report wage patterns. Our goal is to explore the extent to which these data are broadly consistent with the extended version of the model. In summary, we see the model (including the extension incorporating the manual sector) as predicting the following key patterns:

1. During an initial boom phase, we should observe an increase in employment in cognitive and possibly in manual tasks, with an increase in the price of both these tasks. We should witness a decrease in employment in routine tasks and a likely exit out of the home sector (i.e., an increase in the overall employment rate). This process should generally be associated with occupational up-grading for individuals of all skill levels.

2. During the later bust phase, we should observe a decrease in the price of all three tasks. For cognitive and routine tasks we should observe a reduction in employment, while there would be increased employment in manual tasks and a

flow into the home sector. Throughout this phase of the process, we should observe task down-grading for individuals in both high and low education groups.

3.1 Occupational Employment Patterns

In order to document the employment patterns of job categories captured by the model, we group occupations into three broad categories based on the types of tasks predominantly performed within them. The categorization follows that in Acemoglu and Autor (2011) and Autor and Dorn (2013) and consists of: 1) cognitive, non-routine task occupations, including managerial, professional and technical occupations; 2) routine task occupations, which include clerical and office jobs, sales and production occupations; and 3) manual, non-routine task occupations, which include laborers, transportation, farming, and household service occupations. In what follows, we refer to these three occupation groups by their respective predominant task usage.

In Figure 16 we plot the fraction of individuals aged 18-65 employed in occupations that require substantial cognitive skills for each year from 1973 through 2011 (normalized to zero in 1999). As can be seen from the figure, this ratio increased substantially from 1980 to 2000, and then it appears to reach a plateau over the period 2000-2010. On this figure, we also report a (per capita) supply index for cognitive occupations. This index is constructed from a counterfactual exercise. In particular, changes in occupational employment rates could come about through changes in the occupation structure or from composition changes. We perform a simple reweighting exercise that holds the composition of educational attainment, age, and gender constant overtime. Using these weights, we recalculate counterfactual occupational employment rates holding constant the composition of the workforce. We interpret the difference between the observed and counterfactual employment rates as a supply index. This supply measure will reflect changes in, for example, the availability of college graduates for high-skill occupations.

There are two key features of Figure 16 that we wish to highlight. First, from 1980 to about 2000, employment in cognitive jobs grew faster than the supply index, suggesting that demand for cognitive tasks outstripped supply. In contrast, after 2000, the supply index continued to grow at a similar rate as in the pre-2000 period, but cognitive employment stalled. We interpret these trends as suggesting that demand for cognitive jobs likely decreased over this second period since in a simple demand and supply framework, for overall employment to stay constant in the face of increased supply would require a shift down in demand. To make this pattern more transparent, in Figure 17 we adjust our cognitive task employment rate series by subtracting the supply index from the original series, thereby creating a occupational employment rate series adjusted for compositional changes.[19] Superimposed on this series in an estimated cubic trend. As can be seen in the figure, even after adjusting for the change in composition of the population, the ratio of individuals in cognitive jobs increased substantially over the 1980-2000 period. In contrast, after 2000 we see a general fall in this ratio. While the boom period of 2003-2006 reversed the fall for a few years, there appears to be a substantial downward trend post-2000. This constitutes a first piece of evidence suggesting an important reversal in the demand for cognitive tasks beginning in the year 2000.

In Figure 18 we report a series similar to Figure 17 which focuses on routine occupations. This figure controls for changes in the composition of the population in the same way, and overlays a cubic trend. Not surprisingly, the figure shows a fall in the routine-jobs employment rate that starts in the late 1980s, and then gradually accelerates with a very sharp decrease arising in 2000. This pattern has been noticed by many in both the media and academic research (e.g., Charles, Hurst, and Notowidigdo (2012) and Siu and Jaimovich (2012)) and is now commonly interpreted as reflecting the replacement of such jobs by technological advances (robotics and information technology) and off-shoring.

Finally, in Figure 19, we consider manual jobs and obtain a very different pattern from the previous two figures: the adjusted employment rate in manual jobs grew steadily for three decades. This growth arises despite the fact that workforce composition changes have worked against employment in manual jobs. In particular, increases in educational attainment over this period act to reduce supply for manual jobs since these tasks are predominately held by low-skilled workers. Hence,

[19]While we view these exercises as largely indicative rather than conclusive, the issue of the potential endogeneity of educational composition shifts is a clear concern. We have attempted to investigate the importance of this in a number of ways. First, we used state-level data on occupational employment rates and educational composition in regressions of changes in occupational employment rates on changes in the proportion in each educational category. We ran this regression pooled across cities and include year effects, implying that we were using state-specific deviations from overall trends to identify the relationship between education shifts and occupation changes. Using the estimated coefficients from this exercise, we formed predicted occupational employment for the country as a whole and repeat the exercise carried out in Figure 17 using this alternative supply measure. While this uses education coefficients that are constructed very differently from the re-weighting approach, the results are very similar to what we have presented here. There might be concern about using such high frequency variation to identify the education effects and so we also tried using the opposite: a pure between estimator in which we used over-time averages by state. Again, we constructed predicted supply measures using the estimated coefficients. The patterns for the manager, professional and technical occupation group again look very similar to what is presented here.

the sustained growth observed in Figure 19 must reflect a demand for manual tasks that has been strong throughout this time period or that higher-skill workers traditionally supplying their labor to other sectors are now performing manual tasks in response to limited opportunities elsewhere. In what follows, we provide evidence that is supportive of this latter mechanism for the post-2000 period.

Figures 17, 18, and 19 provide an overview of changes in the occupational structure of the working age population when we use a three-group classification to represent jobs. One drawback of such an approach is that such a division is that it is artificial in the sense that occupations generally combine different types of tasks, implying that the classification should be much more continuous. To examine the robustness of the above patterns to a more continuous classification, we adopt a version of Autor, Katz, and Kearney (2006)'s methodology in which we rank occupations by their mean wage in 1980 then group occupations together into 100 categories that correspond to percentiles of employment. This is intended to represent a ranking of occupations by a measure of skill (the average wage). We then calculate the change in the share of total hours worked in the economy for each percentile and fit a smoothed line through these data to get a measure of employment growth for jobs at different percentiles of the occupational skill distribution.[20] Each year of data in the figure actually represents two adjacent years that we pool together to reduce noise. Figure 20 contains the resulting smoothed employment growth changes for the periods 1980-1990, 1990-2000, 2000-2007, 2000-2010.[21]

A few key patterns stand out from this figure. For the decades prior to 2000, employment growth occurs disproportionately in occupations with base-period wages above the median. Given that high-wage occupations tend to be more strongly associated with cognitive tasks and with higher skills, this pattern fits well with theories of skill-biased demand shifts. It is also noteworthy that the changes become non-monotonic in the 1990s, with occupations in the bottom decile of the wage distribution growing relative to those between about the 20th and 40th percentiles. This is the source of Autor, Katz, and Kearney (2006)'s argument that technological change became polarizing in the 1990s in the US.

What is perhaps most striking, though, is the change in the shape of the curve after 2000. The relative growth of the lowest percentile occupations becomes very

[20]A key challenge in constructing these figures is to obtain consistent occupational categories across time. In Appendix A, we provide the details on how we constructed those categories and on robustness check exercises.

[21]Since this figure relates to the share of hours worked in the economy, it is unclear whether the plotted changes correspond to net increases or decreases in employment. In the appendix we therefore report a similar figure calculated for changes in employment rates instead of employment shares.

strong after 2000. This is a point emphasized by (Acemoglu and Autor, 2011) in a figure that is extremely similar to the one presented here. But just as striking is the evaporation of relative growth in the top percentiles. Acemoglu and Autor (2011) note this change in passing but do not emphasize it. We, in contrast, believe this change could be key to understanding the overall shifts emphasized in Section 1.

Since part of the changes in Figure 20 could reflect simple changes in the composition of the working age population, we construct a counterfactual version of Figure 20 in which we hold population composition constant. Our approach is a re-weighting approach similar in spirit to DiNardo, Fortin, and Lemieux (1996), where we form weighted employment totals in each occupation in each year. The weights are constructed so that the re-weighted educational, age, and gender composition in each year remains constant. Specifically, each person's employment outcome is weighted by the ratio of the proportion of people in their education-age-gender category in 1989 to the proportion of people in their category in the survey year. This exercise is intended to effectively un-do occupation specific supply shifts generated from changes in the education composition of the workforce.

Figure 21 contains the plotted employment share change lines using the re-weighted data. The result is an even more strongly U-shaped profile in the 1990s but – and especially relevant for our interpretation of the data – also a decrease in employment share in the 2000s for the occupations at the high end of the wage distribution. We interpret this pattern as a decline in demand for high-wage, cognitive occupations during this decade that was partly masked by an increase in supply of skills to these occupations due to composition changes – particularly the ongoing increases in the educational attainment of the workforce. In both Figures 20 and 21, we report changes for the period 2000-2007 and 2000-2010 in order to clarify the role of the financial crisis, and subsequent adjustments, on the patterns we are emphasizing. As can be seen from either figure, the post-2000 shifts in the occupation structure are not driven by the post-financial crisis period as the data give a similar picture when examining the 2000-2007 or 2000-2010 periods.

In the appendix, we present a double difference taken based on Figure 20, where we plot the differences between the 1980s versus the 1990s and the 1990s versus the 2000s. This serves to highlight that the earlier period from the 1980s to the 1990s was characterized by an acceleration of employment concentration at the top and bottom of the occupational wage distribution. In contrast, for the 2000s versus the 1990s, the major change is a reversal in the employment growth pattern for the high-wage occupations. Thus, it emphasizes what is evident in Figure 20: that there was a 'great reversal' in the demand for cognitive tasks. Later, in the section on task prices, we argue that the wage patterns are supportive of the notion of such a reversal.

3.2 Skill Upgrading and Downgrading

The previous section contrasts the strong growth in cognitive employment prior to 2000 with its subsequent reversal. This pattern pertains to the entire working-age population and is especially striking when controlling for composition shifts in the workforce. In this section, we highlight the response of different education groups to the shift in the demand for cognitive skills. According to the model, workers will respond to an increase in the demand for cognitive skills by moving into these occupations and to a decline in cognitive demand by exiting into alternative sectors. It is these implications we seek to investigate in this section, with the turning point in demand set as 2000.

Figures 22 and 23 examine the task assignment of workers with four years of college education (BAs) using two different measures of task assignment. In Figure 22, we plot an index of the average cognitive task intensity of college graduates over the 1980-2010 period. The index is normalized to equal one in 1990 and movements in the index can be interpreted as changes in average cognitive task usage relative to that year. We measure cognitive intensity by assigning to each 4 digit occupation an average of their scores for cognitive tasks from the Dictionary of Occupation Titles (DOT). We define cognitive tasks as the non-routine analytic and interactive tasks used in Autor, Levy, and Murnane (2003) in their examination of the skill content of jobs.[22] Movements in this cognitive task intensity index reflect movements in college educated workers across occupations. The figure indicates that average cognitive task intensity for college graduates increased from the early 1980s until about the year 2000 and then declined throughout the rest of the series. Figure 23 provides additional evidence of this trend by plotting the ratio of college employment in cognitive (i.e., managerial, professional and technical) jobs to non-cognitive jobs. Similar to the previous figure, the cognitive to non-cognitive employment ratio of college workers increases from the early 1980s until the year 2000. In the 2000s, college workers shifted employment away from cognitive occupations and toward routine and manual occupations.

To further examine the changing patterns of employment among education groups, we construct indices that capture routine and manual task intensities. Our measure

[22] In particular, we measure cognitive tasks by the average of the code `math` and `dcp` provided by David Autor (available on his webpage), which measure 'general educational development: math' and 'accepting the responsibility for the direction, control or planning of an activity.' Precise definitions can be found in Autor, Levy, and Murnane (2003).

for routine task intensity is constructed by averaging Autor, Levy, and Murnane (2003)'s measure for routine-cognitive and routine-manual tasks.[23] Figure 24 plots an index for routine task intensity for both college and high school graduates. As can be seen, the value of this index declines for both education groups throughout the late-1980s and the 1990s, before showing an upward trend for college workers in the 2000s. Following from the model, we interpret the turnaround in the trend for the college educated workers as stemming from them "cascading" down from the cognitive occupations, where demand is falling. For the high school educated workers, after 2000 we see a continued decline in the performance of routine tasks. We interpret this as resulting from the general decline in demand in routine tasks and, after 2000, from more educated workers pushing high school educated workers out of these occupations. It is important to note that there was a substantial change in the occupation coding in the CPS in 2003. This may account for the sharp jump in the proportion of BAs in routine jobs in that year. However, the trend toward the routine sector for the BAs continues to be evident even if we focus only on the post-2003 period where the occupational coding is consistent.

In Figure 25 we present an index of manual task intensity for both college and high school graduates.[24] For high school educated workers, the value of this index rises throughout the time period. In the 1990s, this could reflect the type of increase in demand for these occupations stressed in Autor's work while in our model the increase in the 2000s would reflect a supply shift. For college educated workers, the manual index displays somewhat of a U-shape – declining in the 1980s to the mid-1990s, before trending upward to the end of the series. The reversal in the college series is quite dramatic – college workers become more manual task intensive in the 2000s than at any earlier period in our time frame. [25] This pattern is consistent with the forces emphasized in the model whereby the reduction in the demand for both cognitive and routine tasks in the post-2000 period lead workers from both education groups to supply more work to manual tasks.[26]

To obtain a more detailed view of the occupational shifting by different education groups, Figures 26 and 27 examine changes in employment for young college and high

[23]In particular, we take the average of Autor, Levy, and Murnane (2003)'s variables `sts` and `finger`, which measure 'set standards and tolerances' (routine-cognitive) and 'finger dexterity' (routine-manual), respectively.

[24]The measure for manual task comes form Autor, Levy, and Murnane (2003)'s variable `ehf`, which measures 'eye-hand-foot' coordination.

[25]Again, the 2003 coding shift may be partly responsible for the strong upward pattern but the increase in the manual task index for BAs is evident both before and after 2003.

[26]In the appendix, Figure 36, provides a plot of the cognitive content of jobs held by high school educated workers.

school graduates by comparing their employment densities before and after 2000 for occupations ranked by their average wage in 1980. Occupational densities (or employment rates) are constructed by calculating the probability that an individual in a given education-demographic group is observed in each occupation (including unemployment). Taking the difference in occupational densities over time, when the occupations are ranked by their average wage in 1980, indicates whether a given education group has shifted employment systematically in terms of occupations ranked by wage. In figure 26, we plot these changes for college workers between the ages of 25-35, and fit a line using a local-mean smoother to highlight the pattern. In the first panel, which shows the results for the 1990s, the upward slope of the smoothed line indicates that employment of college graduates shifted towards high-skill, high-pay occupations. The second panel documents that this situation completely reverses for the 2000s: college workers move out of high-wage occupations toward lower-paying ones.[27] We interpret this shift in occupational employment as young college graduates responding to the decline in the demand for cognitive tasks by accepting alternative task assignments. Figure 27 shows that, beginning in the 1990s, high school educated workers were already shifting away from middle-paying occupations towards low-paying ones and that this trend was amplified during the 2000s.

Together, we view the set of Figures from 22 to 27 as providing a consistent picture of shifts in occupational employment for college and high school graduates. In particular, these figures document that over the 1990s there was a shift in employment that is characterized by skill-upgrading of college workers. In the post-2000 period, however, the shifts in employment for both groups of workers can be characterized as one of skill-downgrading, with those shifts having started well before the 2007 financial crisis.

3.3 Task prices

In this subsection, our aim is to provide a picture of the over-time changes in the shadow prices of the three tasks emphasized in the model: cognitive, routine and manual. The behaviour of task prices is key to understanding whether the employment patterns we document are driven by shifts in supply or demand. The narrative that we advance suggests that the decrease in cognitive tasks in the post-2000 period is due to a reduction in demand. If this is the case, we should observe an accompanying decrease in the price of cognitive tasks. On the other hand, we also observe

[27] Appendix Figure 37 shows that the results are similar whether we look over the 2000-2007 or the 2000-2010 period, indicating that this shifting pattern is not the result of the financial crisis but part of an ongoing post-2000 trend.

a shift in employment toward jobs utilizing manual tasks. Our interpretation of the data suggests that this shift is mainly driven by a supply channel; that is, we argue that this shift is the result of high-skilled individuals taking lower skilled jobs due to the decline in demand for both routine and cognitive tasks. As such, this outward shift in supply should place downward pressure on the price of manual tasks.

The difficulty we face in examining these issues is that task prices are not readily observable. We observe the wage paid to workers in different occupations, but, as the model suggests, these wages will not in general reflect the task price. The reason is that the wage paid to an individual employed in a given occupation will reflect the the number of effective units of skill embodied in the individual multiplied by the skill price. As task prices change, so too will the composition of individuals across occupations. The selection mechanism, as parametrized in the model, implies that changes in the price of the cognitive task would not be reflected in changes in the average wage in cognitive jobs. While we do not take this parametrization seriously, we do believe that selection is relevant over this period, and this makes inferences about task prices from observed wage movements difficult.

To illustrate the potential importance of selection of individuals across occupations over time, in Figure 28 to 30 we plot two alternative measures of the average wage paid in cognitive, routine and manual jobs. The first measure, represented by the dark line in each of the figures, corresponds to the simple, observed average wage paid in each occupation. The second measure, represented by the dashed line in each figure, calculates the average wage in each occupation while holding the composition of education, age and gender constant at their 1980 levels. When focusing on the 2000-2010 period, the average wage for each of the three occupation groups increases substantially. From this perspective, it would not appear that the prices for any of these tasks had declined. However, when we control for changes in the observable characteristics of the individuals in each of these occupations, we get a very different picture. Composition adjusted changes in real wages over the period are close to zero for each group. For example, the real wage growth over 2000-2010 for wages in cognitive occupations is 6% in the raw data and about 2% when adjusted for observable. Similarly, for manual tasks the growth in real wages is close to 6% in the raw data and close to 1% in the adjusted data over the same period. Since this correction only accounts for changes in observables, it likely underestimates the effects of selection on the wage series as it is quite plausible that changes in unobserved heterogeneity mimicked changes in observed heterogeneity.

We examine the relevance of selection on unobserved heterogeneity in Figure 31, where we plot the changes in real wages of young college educated workers at each percentile of the wage distribution. In constructing this figure, we calculate each percentile of the wage distribution in each year. In doing so, we impute the wages of non-workers as zero, which accounts for about a 5th of the sample for each year on average. The series in the figure represent the differences in log wages at each percentile of non-zero wages for the indicated time periods.[28] To relate these changes to task prices (which is our ultimate goal), requires an assumption of a 'single-index' model of wage generation. Given this, if the distribution of the unobserved skill is relatively constant over time then we can interpret changes in wages at a given percentile as representing the change in the wage for a worker with a given skill level defined by a combination of observed and unobserved skills.[29] If, in addition, the workers at a given percentile are predominantly in one occupation then we can go further and interpret the wage change at that percentile as corresponding to the change in the task price (the amount paid per unit skill) in that occupation. We recognize that the assumptions underlying this exercise are strong but, in the absence of true panel data, we still view them as reasonable (though crude) approximations to the skill prices that are our main point of interest.

As can be seen from the figure, over the 1990-2000 period wages increased substantially for young college workers in the top quarter of the wage distribution. Since cognitive jobs are heavily represented in this part of the wage distribution (with approximately 80% of young BA workers with wages between the 80th and 100th percentiles being in cogntive occupations), we interpret this increase as reflecting an increase in the price of cognitive tasks during this time. When we turn to the post-2000 time period, a different picture emerges: real wages decrease throughout the distribution for young college workers. Notably, this includes wages at the top of the distribution, which, again, mainly correspond to employment in cognitive jobs, implying that the price of cognitive tasks has likely fallen over this period.[30]

Figure 32 constructs a similar figure for young high school graduates. As is clear from the figure, the proportion of this group not working is much larger and this results in the change in wage percentile series beginning to take non-zero values

[28]For percentiles at which individuals are non-employed in one year but employed and with an observed wage in another year, we plot the difference as a zero.

[29]Note that it is important to include the non-workers in this exercise because the 90th percentile of the wage distribution conditional on working will shift as the employment rate shifts, implying that it would correspond to workers with different underlying skill index values.

[30]Recall that, in our model, the entrepreneur group will benefit throughout. Some of these individuals may be in our sample which would tend to bias upwards the real wage movements and thereby could mask any fall in the price of cognitive tasks to some extent. For older (over age 35) BAs, where we believe this problem is likely more prevalent (in addition to potential biased due to more long term implicit contracts), we do not observe declines in real wages at the very top of the wage distribution. This is reported in Figure 39 of the appendix.

at much higher percentiles than for the college group. Focusing on the 1990-2000 period, the pattern of wage changes is 'J' shaped, with increases at the bottom of the wage distribution while being stagnant or declining at the top. For young high school workers, those with wages between about the 80th and 90th percentiles are predominantly (over 60%) in routine jobs. Focusing on this part of the distribution indicates a slight fall in real wages, which we interpret as a fall in the routine task price in this period. At the other end, young high school workers below about the 45th percentile in this figure are largely (close to 60%) in manual occupations. In that range, there is an increase in wages in the 1990s, which we interpret as a rise in the manual task price. Taking Figures 31 and 32 together, the patterns we document are consistent with those emphasized by Autor, Katz, and Kearney (2008) who view the 1990s as a period of wage 'polarization' brought about by changes in task demands. Turing to the 2000s, the figure indicates that young high school workers experienced a fall in wages at all percentiles. The decline in the region we have identified with routine tasks is about the same size as in the 1990s but the pattern near the 40th percentile, which we associate with the manual task price, is strongly downward. The sharp fall in wages at the bottom fits with our narrative that the increase in employment in manual jobs reflects supply factors rather than demand factors for the post-2000 period.

To bring together these implications in a more readable manner, in Figure 33, we plot our measures of the task prices for the three occupations for each year in our sample. More specifically, the cognitive task price corresponds to the average of the log wages between the 80th and 90th percentiles of the distribution (including the non-employed) for young BAs. The routine task price corresponds to the average of the log wage between the 80th and 90th percentiles for young high school individuals, and the manual task price is the average between the 37th and 40th percentiles for the young high school individuals. We normalize all three series to their 1990 values. From this one can see that the early 90s recession was associated with declines in all three task prices, offset to some degree by gains in the strong labor market of the second half of the 90s. For cognitive tasks, the gains were strong enough that the task price was above its 1990 price by the early 2000s, and the same is true to a lesser extent for manual tasks. In contrast, the routine task price does not recover its 1990 level. Again, this fits with the standard story of polarization in the 1990s. But all three task prices fall in the 2000s and, notably, are declining well before the onset of the late 2000s recession. Our measure of the cognitive task price falls by about 2% in real terms, fitting with our story of declining demand for these tasks. This decline is dwarfed by the 8% decline in the price of manual tasks. The latter drop fits with our claim that the manual task market in the 2000s can be characterized as being dominated by an outward supply shift.[31]

To complete the wage picture, in Figure 34 we report changes in real wage by deciles for young workers when we do not divide by education level. This has both advantages and disadvantages relative to the earlier figures controlling for education. On the one hand, if there has been sorting of individuals between education groups over the period in question, the distribution of unobservable skill within an education group may change, and comparing percentiles while conditioning on education is less clean as a measure of task prices. On the other hand, increases in educational attainment over time may change the quality at a given percentile. Hence, we choose to present results with and without conditioning on education. When we do not condition on education, as in Figure 34, the 1990s show a clear 'U' shaped wage growth pattern that has been documented elsewhere (see, for example, Autor, Katz, and Kearney (2008) and Lemieux (2008)) and has become known as 'wage polarization.' The 2000s, on the other hand, show a fall in real wages at all percentiles with the greatest declines at the low end of the wage distribution.

3.3.1 Cross-City Evidence

We next look for corroborating evidence for the model by using very different data variation from what we have used to this point: cross-city variation. In particular, we treat each city as a separate economy with differential abilities to take advantage of the new technology. In the context of the model, this corresponds to assuming that local economies can differ in their change in ϕ, the technology improvement. Differences in the capacity of cites to take advantage of new technological opportunities could arise from difference in local culture or local advantages, but we will remain agnostic about the specific reason for these capacity differences.

Given these differences across local economies, cities with a greater change in ϕ (i.e., cities better situated to take advantage of the technological shift) should have seen a greater boom in the 1980 and 1990s, especially in terms of increased hiring in managerial and professional occupations. However, such cities should also experience

[31] The post-2000 pattern of declines does not perfectly fit with our Roy type model of supply. In particular, for workers to be choosing to move to lower occupations as the model implies, we should see the decline in task prices being largest for the high end occupations and smallest for the low end, service occupations. One could view the relative movements in the cognitive and routine task prices as potentially in line with this since they both move down to similar degrees and there surely is considerable measurement error in our task price measures. But the much larger decline in the manual task price is hard to reconcile with this feature of the model. As discussed earlier, such a pattern might arise if the routine occupations face downward wage rigidities arising, for example, because of unions.

a greater bust post-2000. To explore this implication – which suggests that the post-2000 decline reflects an ongoing adjustment to the earlier boom – we examine the extent to which the post-2000 bust is associated with measures of the pre-2000 boom. In particular, our measure of the bust is the change in the employment rate. We begin by focusing on changes in the aggregate employment rate, but we also provide evidence broken down by education groups.

Given our assumption that adoption of the new technology requires managerial workers, higher adoption cities should be cities with more substantial increases in managers both as a proportion of the population of the city and as a share of employment in the city in the 1990s. The higher-adoption cities should also see a specific set of patterns in the 2000s that can be explored using a common regression framework as follows:

$$\Delta ER_{c,2010-2000} = \alpha_0 + \alpha_1 X_{c,pre-2000} + \epsilon_{ct}, \tag{8}$$

where $\Delta ER_{c,2010-2000}$ is the change in the employment rate in city c post-2000, $X_{c,pre-2000}$ is an indicator of the boom in the pre-2000 period, and ϵ_{ct} is an error term. In addition, α_0 is a constant, which allows for a common trend for all cities. Our main focus is on the sign of α_1, which should be negative if our boom-bust interpretation is valid. In particular, we consider several indicators of the boom: the change in the employment rate of managers, either over the 1980-2000 period or over the 1990-2000 period; the growth in managerial employment over these two periods; and the simple change in the employment rate over 1980-2000 and 1990-2000.

Table 2 presents the results of OLS estimates of equation (8). The data we use in all of our city-level cross section estimations comes from the U.S. Census and ACS data from 1980-2010. The dependent variable in all estimates contained in Table 2 is the change in employment rate for both men and women from 2010-2000. Each column contains the estimate of α_1 obtained using a specific measure of $X_{c,pre-2000}$. In particular, in columns 1 and 2, the regressors are the change in the employment rate of Managers, Professionals, and Technicians for the 1990s and for the whole 1980-2000 period, respectively. Columns 3 and 4 contain the estimates when we use the growth rates of those same employment rates, and columns 5 and 6 report estimates using the change in the overall city employment rate. In all the columns, estimates of the α_1 coefficient are negative and statistically significant. While these results are just simple associations, they are supportive of the boom-bust interpretation of the model.

3.4 How important could the reversal in cognitive skill demand be in explaining the current low rate of employment?

In this subsection, we wish to quantify the importance of the forces we have identified (the reversal in demand for cognitive tasks and the cascading down of supply after 2000) for understanding current low rates of employment in the US. The precise counterfactual to consider in this exercise is not easy to discern. We address this issue by asking an extremely simple question: How much higher would employment after 2000 have been if:

1. The growth in demand for cognitive tasks had been as great as in the pre-2000 period,
2. All workers displaced from cognitive occupations directly push out workers in other sectors, one for one
3. The greater increase in the demand for cognitive jobs that would have occurred if the pre-2000 trend had continued would not have increased or decreased the demand for routine or manual tasks.

Under this extreme scenario, which we view as a clear upper bound on the potential effects of the the reversal of cognitive demand, we can simply use Figure 16 and project the trend growth in the pre-2000 period in cognitive employment to the post-2000 period and take the difference relative to the actual outcome.[32] Doing this simple calculation we find that the employment rate would be about 5% higher today. While we do not claim that this counterfactual is very meaningful, we do believe that it highlights (as an upper bound) the potential quantitative importance of the reversal in skill demand in adding to our understanding of the fall in employment rates since 2000.

4 Conclusion

As we noted at the outset, a substantial disagreement exists about the causes behind the current low rate of employment in the US. Cyclical effects of the 2008 financial

[32]Since the trend growth in supply for cognitive jobs is rather similar across the pre and post 2000 period, using either Figure 4 or 5 to do this exercise gives a similar answer. In the calculation presented here we use Figure 4 as the estimation of a trend is easier.

crisis likely play a role, and the structural decline in employment in routine occupations and manufacturing jobs are certainly contributing factors (Charles, Hurst, and Notowidigdo, 2012; Siu and Jaimovich, 2012). In this paper, we present theory and evidence suggesting that to understand the current low rates of employment in the US one needs to recognize the large reversal in the demand for skill and cognitive tasks that took place around the year 2000. In particular, we have argued that after two decades of growth in the demand for occupations high in cognitive tasks, the US economy reversed and experienced a decline in the demand for such skills. The demand for cognitive tasks was to a large extent the motor of the US labor market prior to 2000. Once this motor reversed, the employment rate in the US economy started to contract. As we have emphasized, while this demand for cognitive tasks directly effects mainly high skilled workers, we have provided evidence that it has indirectly affected lower skill workers by pushing them out of jobs that have been taken up by higher skilled worker displaced from cognitive occupations. This has resulted in high growth in employment in low skilled manual jobs with declining wages in those occupations, and has pushed many low skill individual;s out of the labor market.

To help organize our thoughts about this process, we presented a simple model where both the pre-2000 boom and post-2000 bust in demand for cognitive tasks could be interpreted as the result of one underlying force in the form of the diffusion of skilled-biased technological change. The only difference with more conventional models of skill-biased technological change is our modelling of the fruits of cognitive employment as creating a stock instead of a pure flow. This slight change causes technological change to generate a boom and bust cycle as is common in most investment models. We also incorporated into this model a standard selection process whereby individuals sort into occupations based on their comparative advantage. The selection process is the key mechanism that explains why a reduction in the demand for cognitive tasks, which are predominantly filled by higher educated workers, can result in a loss of employment concentrated among lower educated workers. While we do not claim that our model is the only structure that can explain the observations we present, we believe it gives a very simple and intuitive explanation to the changes pre- and post-2000.

References

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

Acemoglu, D., and D. H. Autor (2011): "Skills, tasks and technologies: Implications for employment and earnings," *Handbook of Labor Economics*.

Autor, D., F. Levy, and R. Murnane (2003): "The Skill Content of Recent Technological Change: An empirical exploration," *Quarterly journal of economics*, 118(4), 1279–1333.

Autor, D. H., and D. Dorn (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103(5), 1553–1597.

Autor, D. H., L. F. Katz, and M. S. Kearney (2006): "The Polarization of the Labour Market," Discussion Paper January, NBER.

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

Beaudry, P., and D. A. Green (2005): "Changes in U. S. Wages, 1976-2000: Ongoing Skill Bias or Major Technological Change?," *Journal of Labor Eocnomics*, 23(3), 1976–2000.

Carlaw, K., and R. Lipsey (2002): "Externalities, technological complementarities and sustained economic growth," *Research Policy*, 31, 1305–1315.

Charles, K., E. Hurst, and M. Notowidigdo (2012): "Manufacturing Busts, Housing Booms, and Declining Employment: A Structural Explanation," Discussion paper.

DiNardo, J., N. M. Fortin, and T. Lemieux (1996): "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 64(5), 1001–1044.

Dustmann, C., J. Ludsteck, and U. Schönberg (2009): "Revisiting the German Wage Structure," *Quarterly Journal of Economics*, 124(2), 843–881.

FIRPO, S., N. M. FORTIN, AND T. LEMIEUX (2011): "Occupational Tasks and Changes in the Wage Structure Occupational Tasks and Changes in the Wage Structure," Discussion Paper 5542, IZA.

GOOS, M., AND A. MANNING (2007): "Lousy and lovely jobs: The rising polarization of work in Britain," *The Review of Economics and Statistics*, 89(February), 118–133.

JAEGER, D. A., AND M. E. PAGE (1996): "Degrees Matter: New Evidence on Sheepskin Effects in the Returns to Education," *The Review of Economics and Statistics*, 78(4), 733–740.

JUHN, C. (1999): "Wage inequality and demand for skill: evidence from five decades," *Industrial and Labor Relations Review*, 52(3), 424–443.

KATZ, L. F., AND K. M. MURPHY (1992): "Changes in relative wages, 1963-1987: supply and demand factors," *The Quarterly Journal of Economics*, 107(1), 35–78.

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

——— (2008): "The changing nature of wage inequality," *Journal of Population Economics*, 21(1), 21–48.

PARK, J. H. (1994): "Estimation of Sheepskin Effects and Returns to Schooling Using he Old and the New CPS Measures of Educational Attainment," Discussion Paper 717, Princeton University, Industrial Relations Section.

SIU, H. E., AND N. JAIMOVICH (2012): "The Trend is the Cycle: Job Polarization and Jobless Recoveries," Discussion paper, NBER Working Paper No. 18334.

Figure 1:

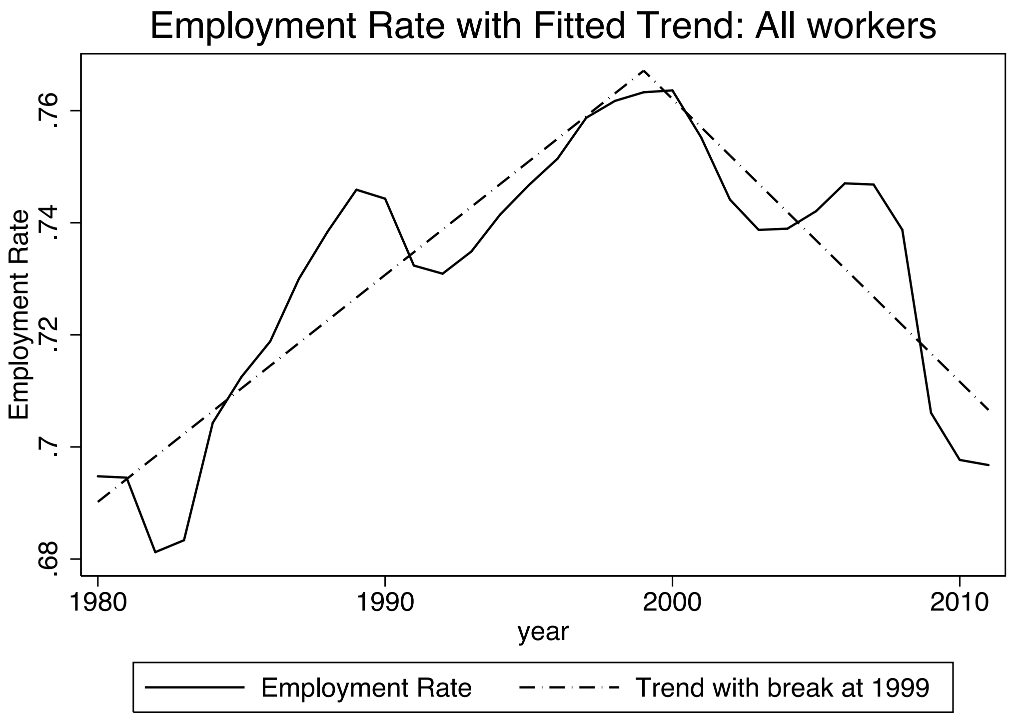

notes: The figure uses ORG data from 1980 to 2011. The employment rate is calculated by summing the number of respondents employed during the survey week over the total population aged 18-64 with positive potential work experience. Mean hourly log wages calculations exclude allocated wages.

Figure 2:

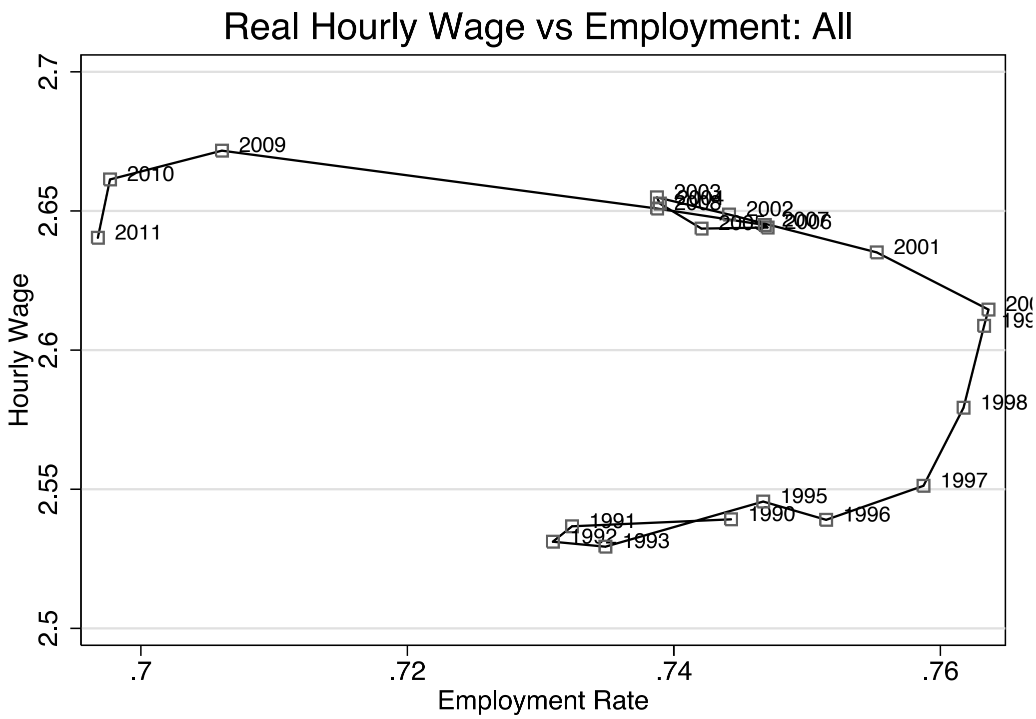

notes: The figure uses ORG data from 1980 to 2011. The employment rate is calculated by summing the number of respondents employed during the survey week over the total population aged 18-64 with positive potential work experience. Mean hourly log wages calculations exclude allocated wages.

Figure 3:

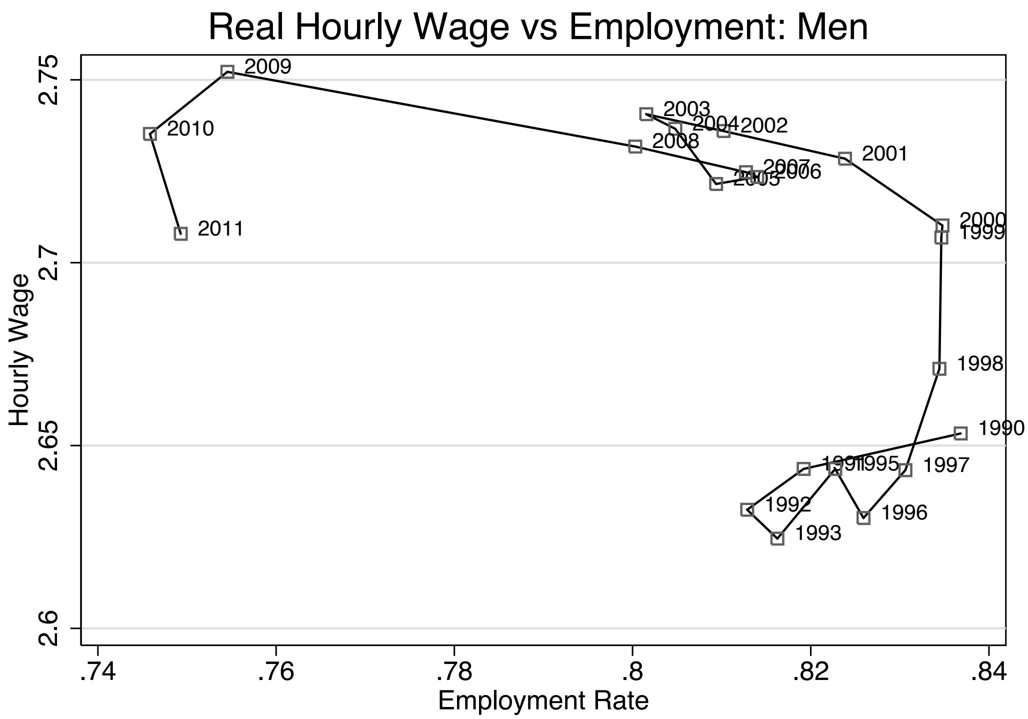

notes: The figure uses ORG data from 1980 to 2011. The employment rate is calculated by summing the number of respondents employed during the survey week over the total population aged 18-64 with positive potential work experience. Mean hourly log wages calculations exclude allocated wages.

Figure 4:

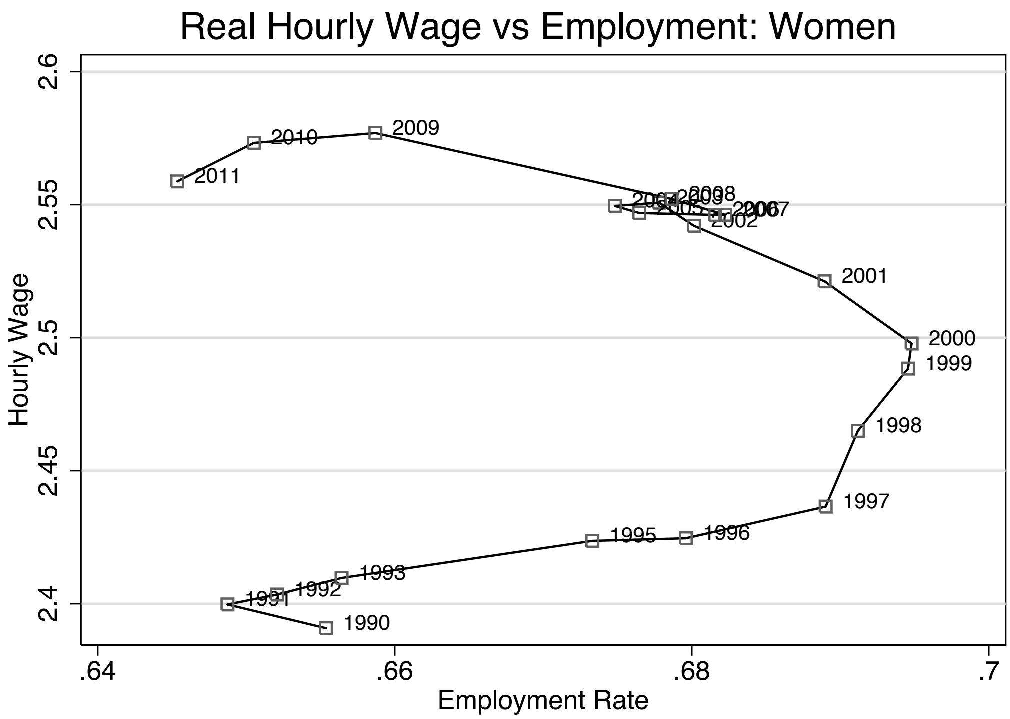

notes: The figure uses ORG data from 1980 to 2011. The employment rate is calculated by summing the number of respondents employed during the survey week over the total population aged 17-64 with positive potential work experience. Mean hourly log wages calculations exclude allocated wages.

Figure 5:

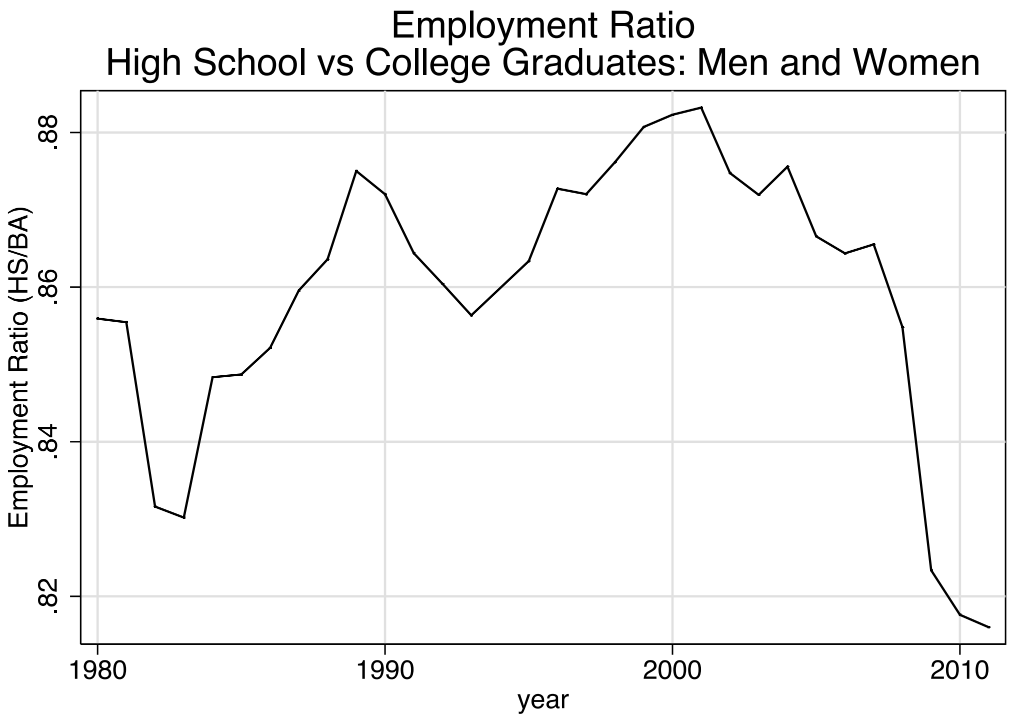

notes: The figure uses ORG data from 1980 to 2011. Employment rates are calculated for workers with exactly 16 or 12 years of education between the ages of 25-65.

Figure 6:


Investment as a share of GDP
Equipment and Software
10
9
8
7
6
5
1950
1960
1970
1980
1990
2000
2010
Quarter

Figure 7:


Investment as a share of GDP
Information Processing Equipment and Software
0
1
2
3
4
5
1950
1960
1970
1980
1990
2000
2010
Quarter

Figure 8:

Investment as a share of GDP
Computers and Peripheral Equipment
0
.2
.4
.6
.8
1
1950
1960
1970
1980
1990
2000
2010
Quarter

Figure 9:

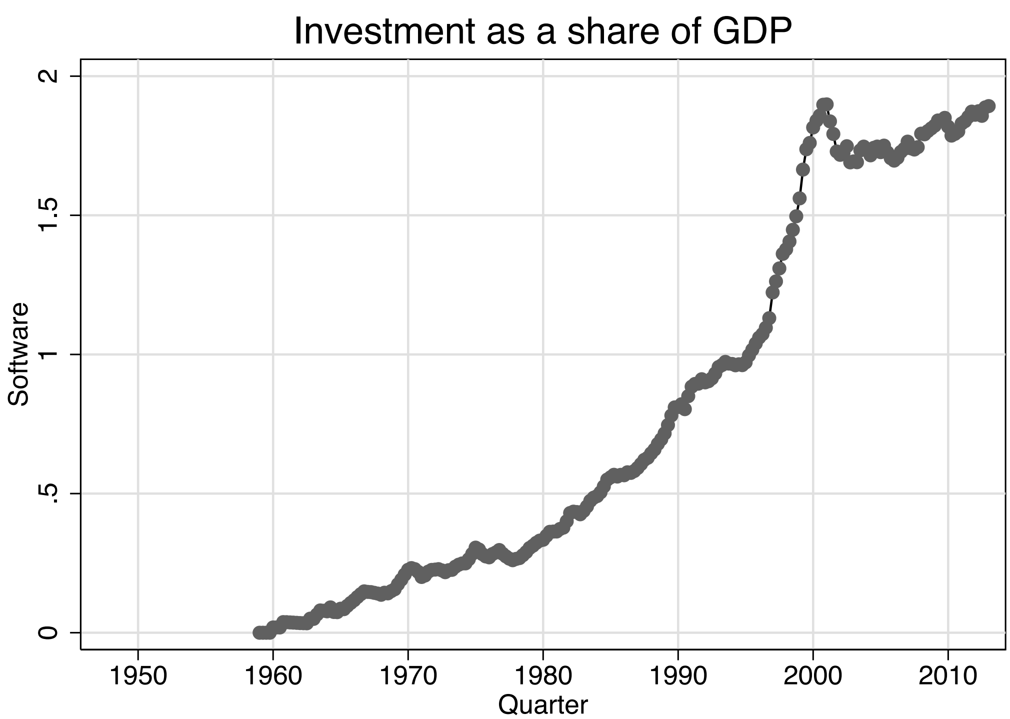

Investment as a share of GDP
Software
0
.5
1
1.5
2
1950
1960
1970
1980
1990
2000
2010
Quarter

Figure 10:

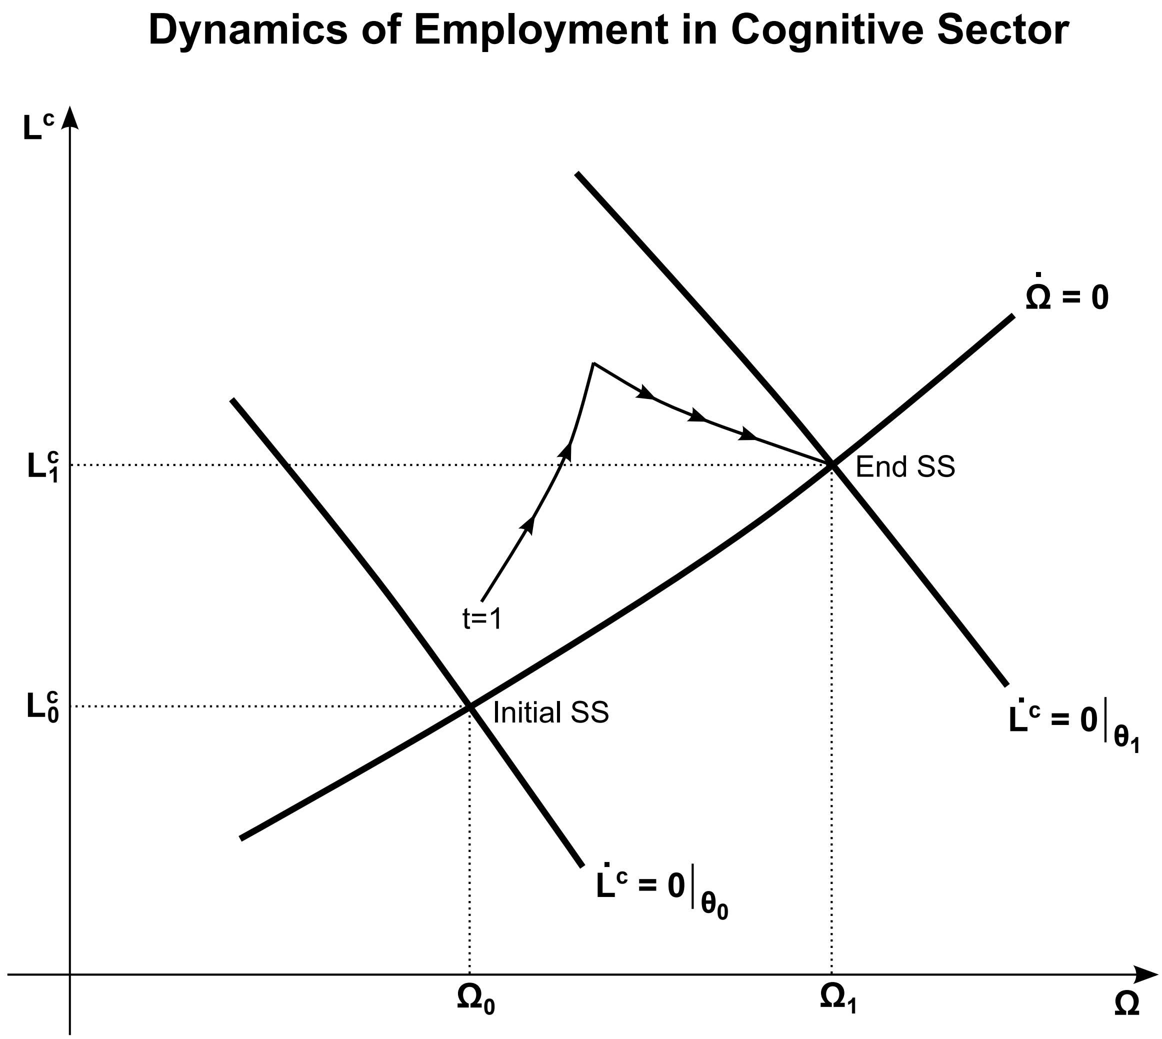
Dynamics of Employment in Cognitive Sector
L^c
$\dot{\Omega} = 0$
L^c_1
End SS
t=1
L^c_0
Initial SS
$\dot{L}^c = 0\big|_{\theta_1}$
$\dot{L}^c = 0\big|_{\theta_0}$
Ω_0
Ω_1
Ω

Figure 11:

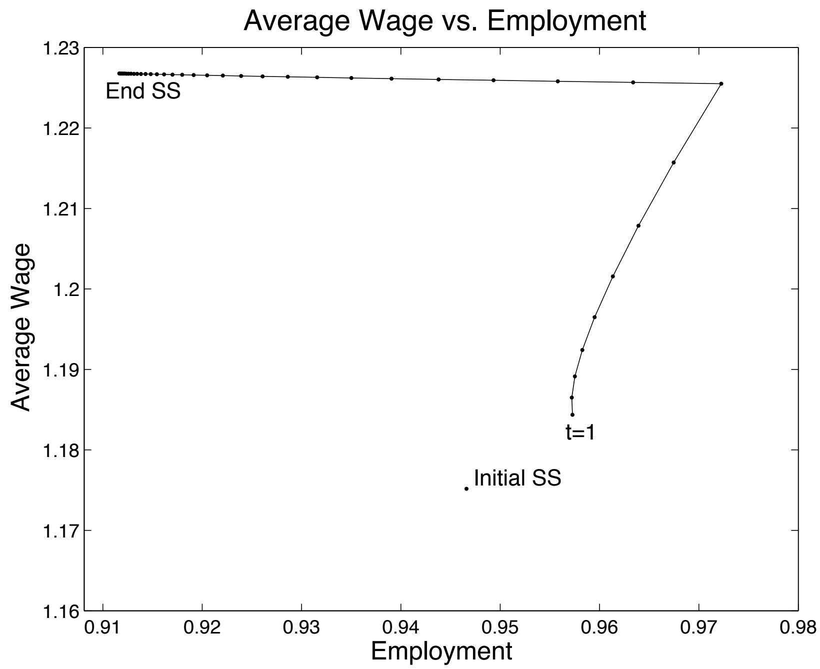

notes: This figure plots the response of an economy to a shock in θ, the technology parameter. In particular, the plot shows an economy in steady state at time $t = 0$ and learns at $t = 1$ of an improvement in θ at $t = 10$. The details of the parametrization can be found in the text.

Figure 12:

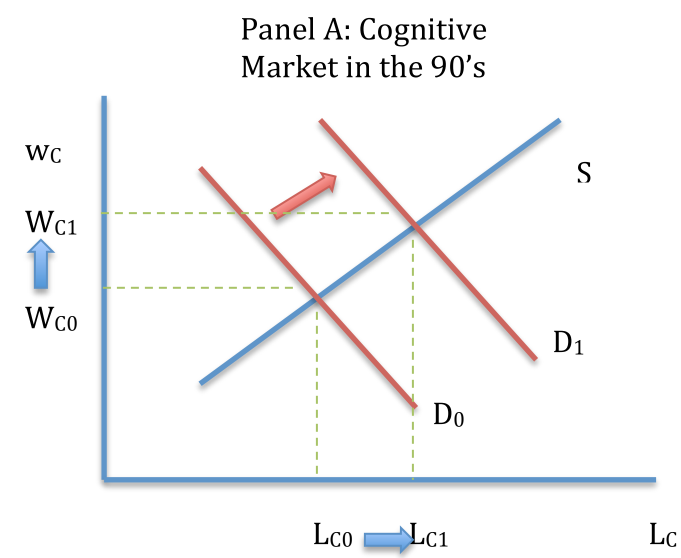

Panel A: Cognitive
Market in the 90's
wC
WC1
WC0
S
D1
D0
LC0
LC1
LC

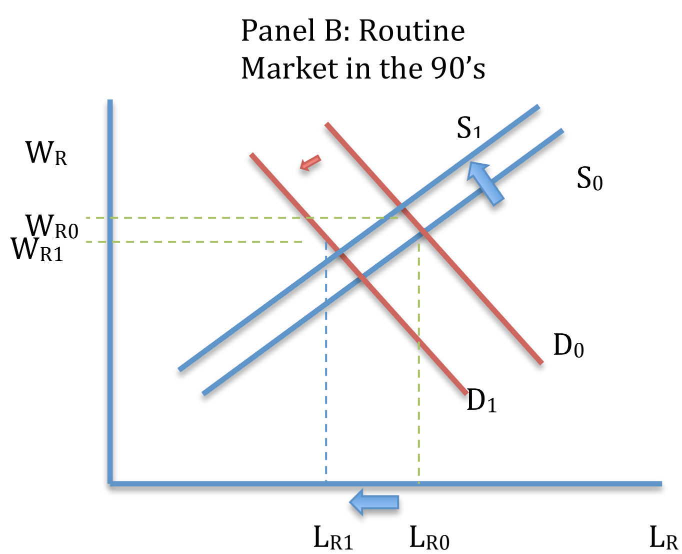

Panel B: Routine
Market in the 90's
WR
WR0
WR1
S1
S0
D0
D1
LR1
LR0
LR

Figure 13:

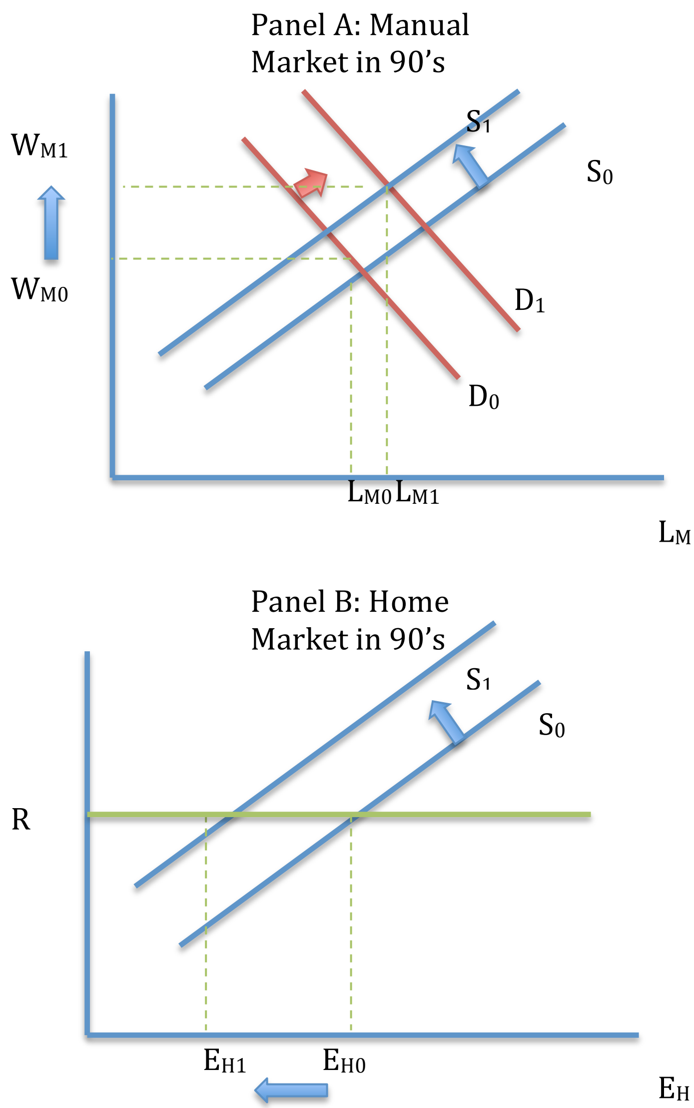
Panel A: Manual Market in 90's
W_{M1}
W_{M0}
S_1
S_0
D_1
D_0
L_{M0} L_{M1}
L_M
Panel B: Home Market in 90's
S_1
S_0
R
E_{H1}
E_{H0}
E_H

Figure 14:

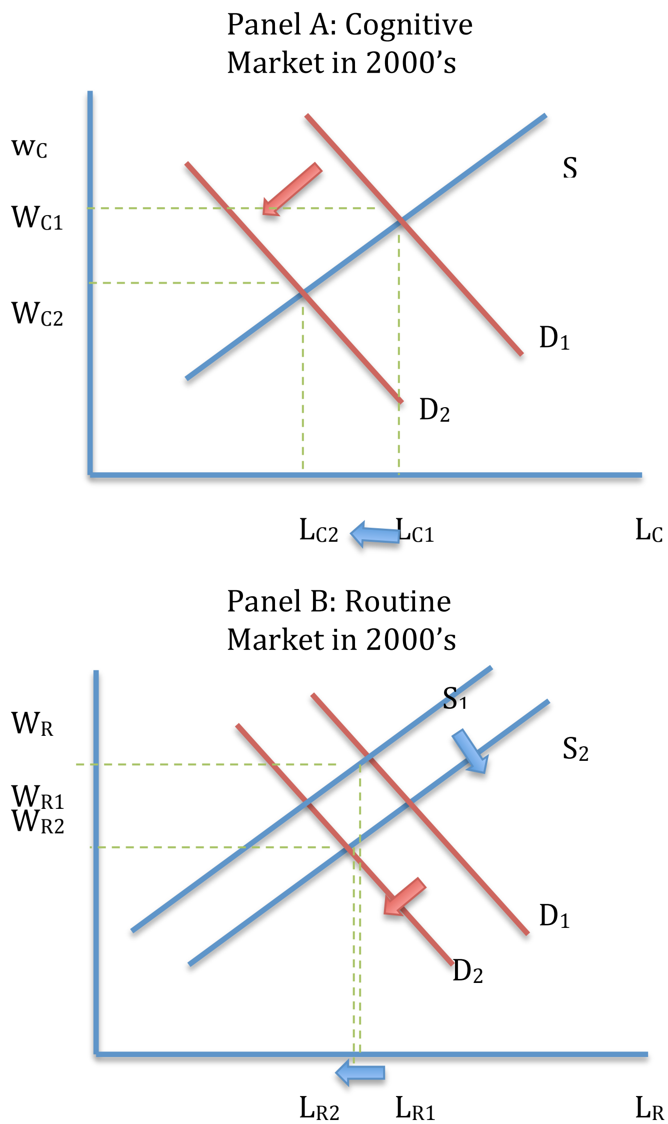
Panel A: Cognitive Market in 2000's
wC
WC1
WC2
S
D1
D2
LC2
LC1
LC
Panel B: Routine Market in 2000's
WR
WR1
WR2
S1
S2
D1
D2
LR2
LR1
LR

Figure 15:

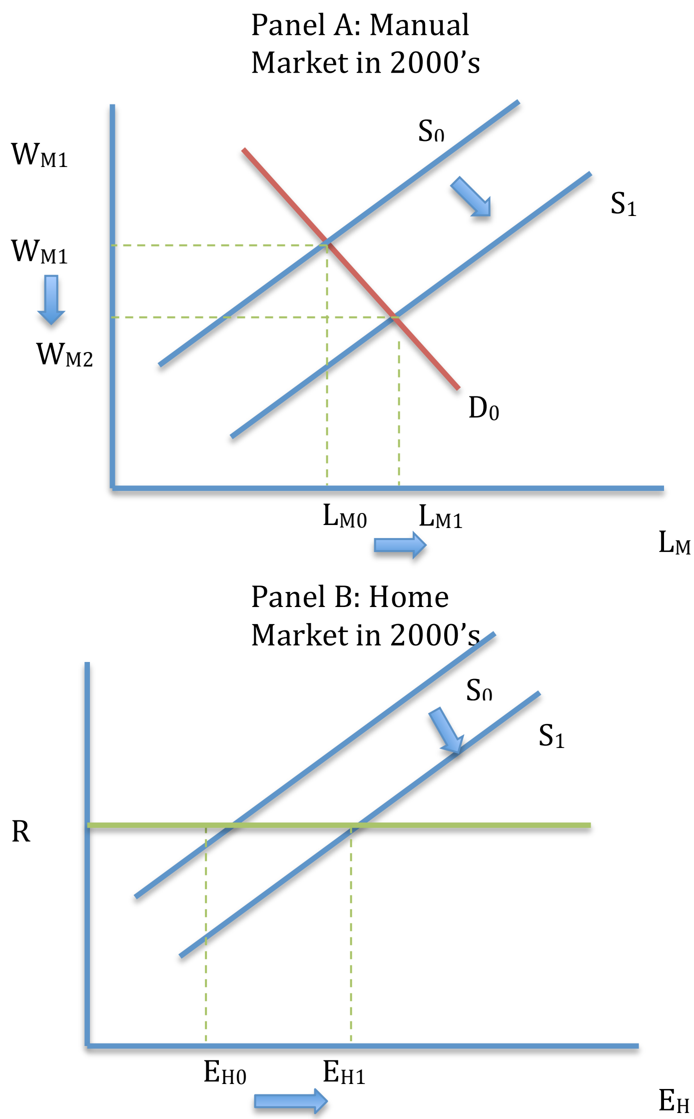
Panel A: Manual
Market in 2000's
W_{M1}
S_0
S_1
W_{M1}
W_{M2}
D_0
L_{M0}
L_{M1}
L_M
Panel B: Home
Market in 2000's
S_0
S_1
R
E_{H0}
E_{H1}
E_H

Figure 16:

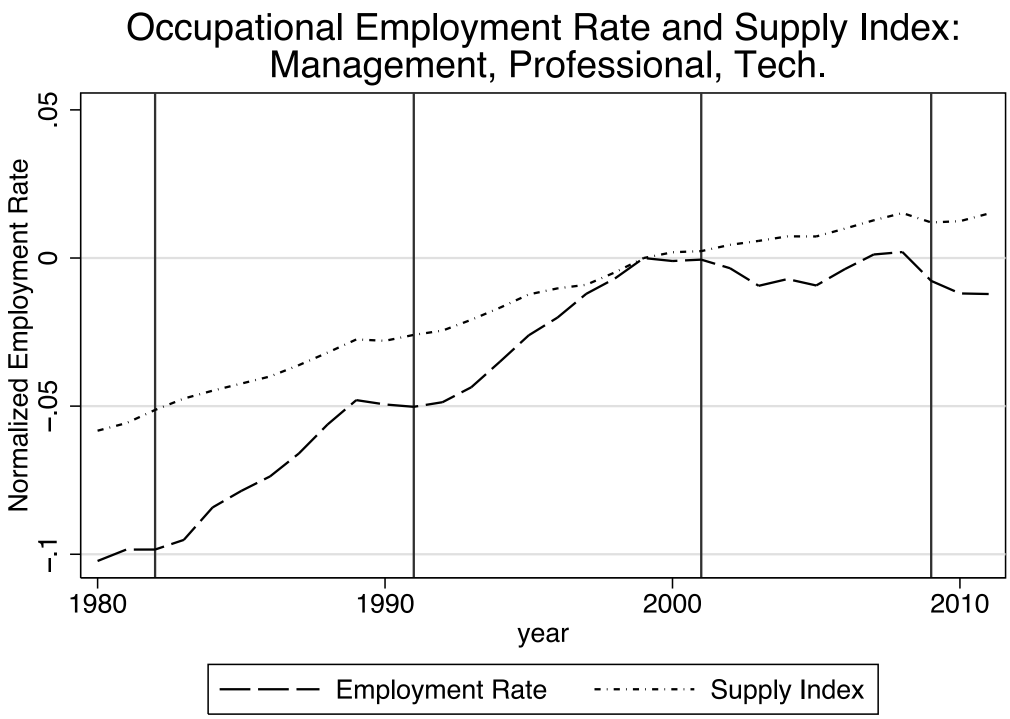

notes: The figure uses ORG data from 1980 to 2011. The employment rate in the figure is calculated as total hours worked in cognitive jobs over the size of the workforce. The supply index is constructed as described in text. Both measures are normalized to 0 in 1999.

Figure 17:

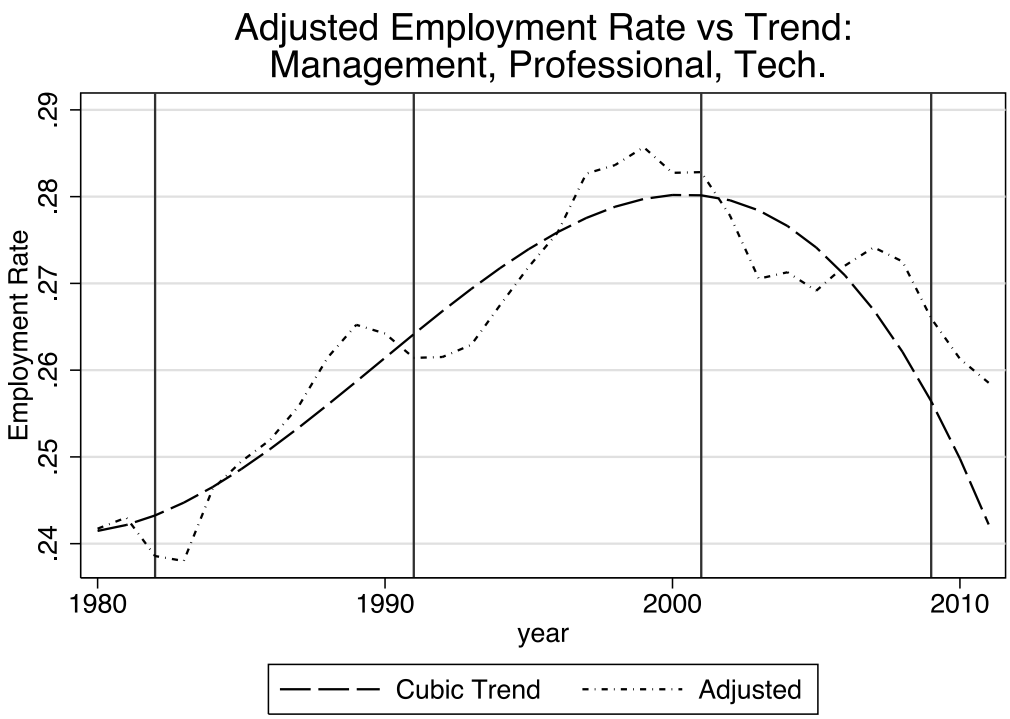

notes: The data in the figure comes from CPS ORG from 1980-2011. The figure plots the employment rate in Cognitive jobs after adjusting for composition shifts, as described in the text, and a fitted cubic trend.

Figure 18:

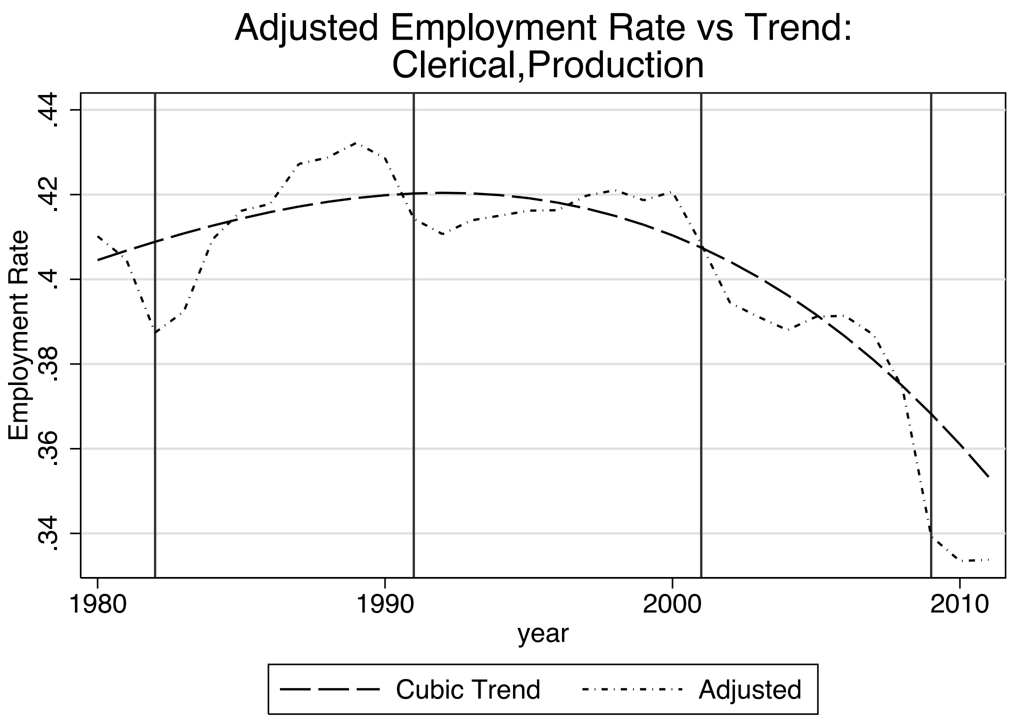

notes: The data in the figure comes from CPS ORG from 1980-2011. The figure plots the employment rate in Routine jobs after adjusting for composition shifts, as described in the text, and a fitted cubic trend.

Figure 19:

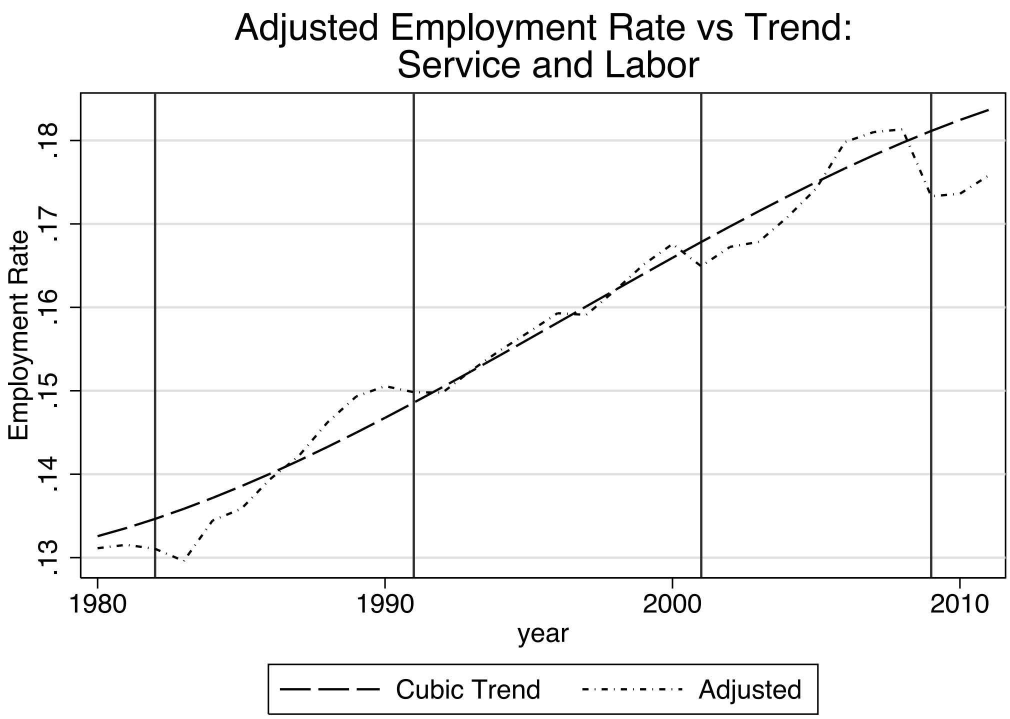

notes: The data in the figure comes from CPS ORG from 1980-2011. The figure plots the employment rate in Manual jobs after adjusting for composition shifts, as described in the text, and a fitted cubic trend.

Figure 20:

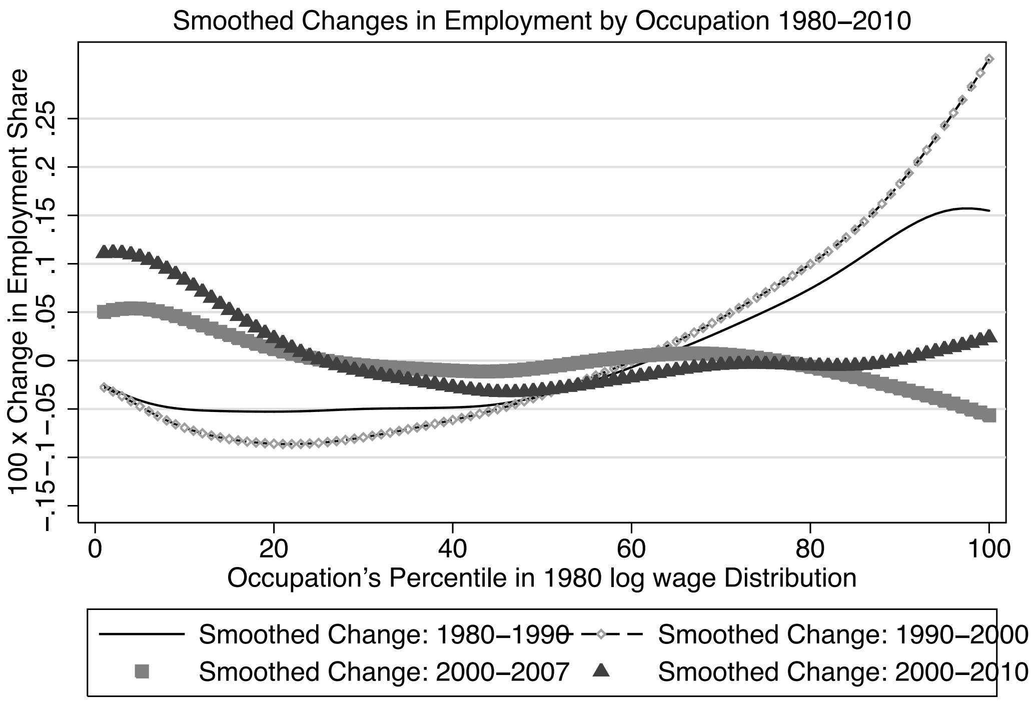

notes: The figure plots the log changes in employment shares by 1980 occupational log wage percentile rank using a locally weighted smoothing regression (STATA's `lowess` smoother with bandwidth 0.8 with 100 observations). Employment shares refer to shares of hours worked in the economy. Occupation codes used in the figure are based on the 1980 CPS occupational coding scheme made consistent across occupational code breaks in 1983 and 2003 using the cross-walks weights provided by BLS. Further details are provided in the data appendix.

Figure 21:

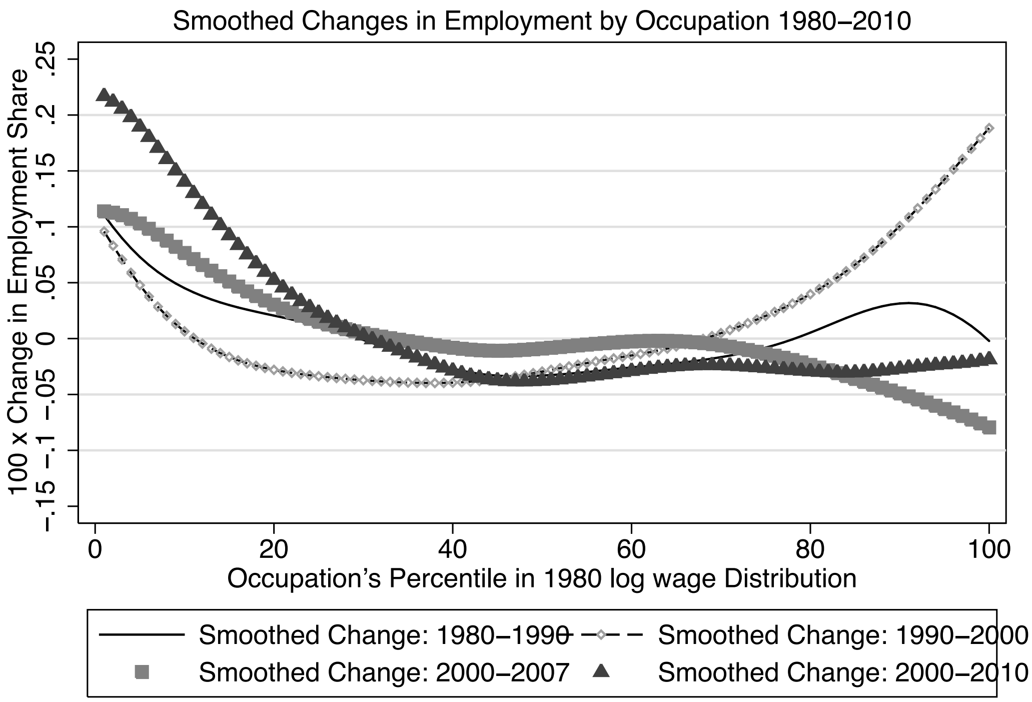

notes: The figure plots the log changes in reweighted employment shares by 1980 occupational log wage percentile rank using a locally weighted smoothing regression (STATA's `lowess` smoother with bandwidth 0.8 with 100 observations). The reweighted employment shares are calculated by holding the demographic characteristics of the population constant as described in text. Employment shares refer to shares of hours worked in the economy. Occupation codes used in the figure are based on the 1980 CPS occupational coding scheme made consistent across occupational code breaks in 1983 and 2003 using the cross-walks weights provided by BLS. Further details are provided in the data appendix.

Figure 22:

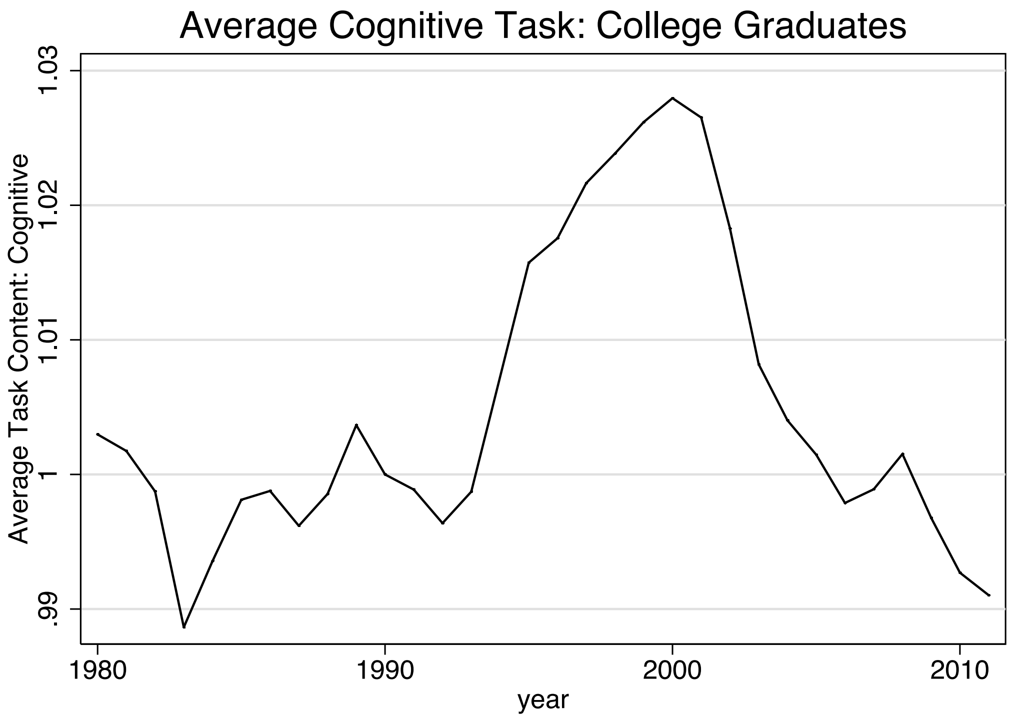

notes: The figure plots an index (normalized to one in 1990) of average level of cognitive task for employed college graduates over time using the CPS ORG data from 1980-2011. The cognitive task measure comes from the average of the variables `math` and dcp described in Autor, Levy, and Murnane (2003) and in the data appendix.

Figure 23:

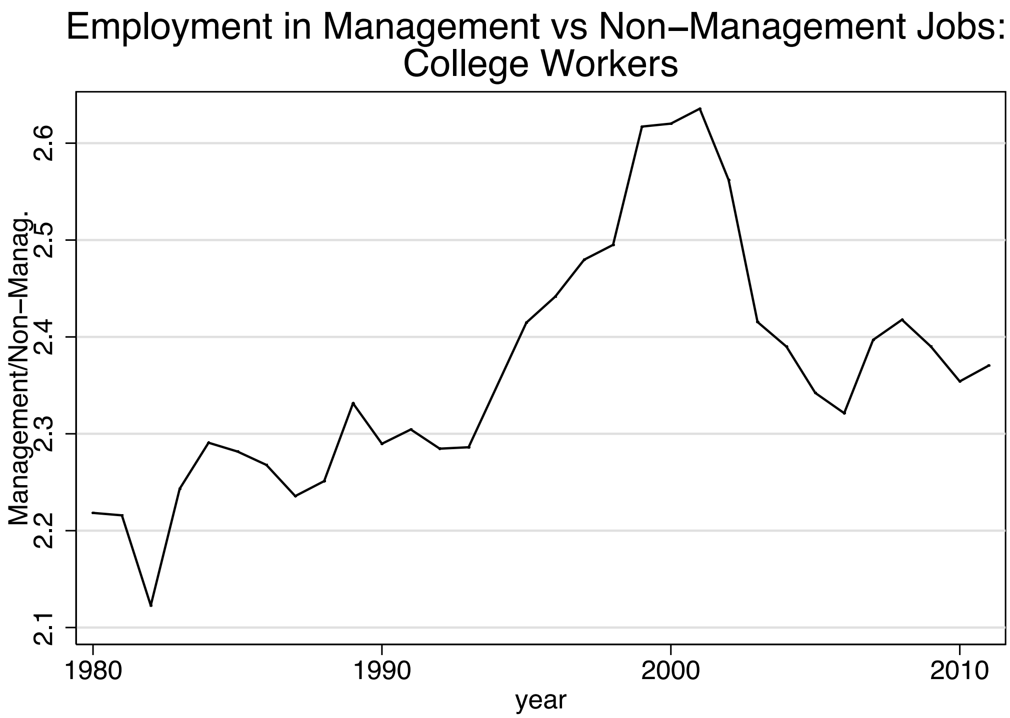

notes: The figure plots the ratio of employment in cognitive jobs vs non-cognitive jobs for college graduates over time using the CPS ORG data from 1980-2011.

Figure 24:

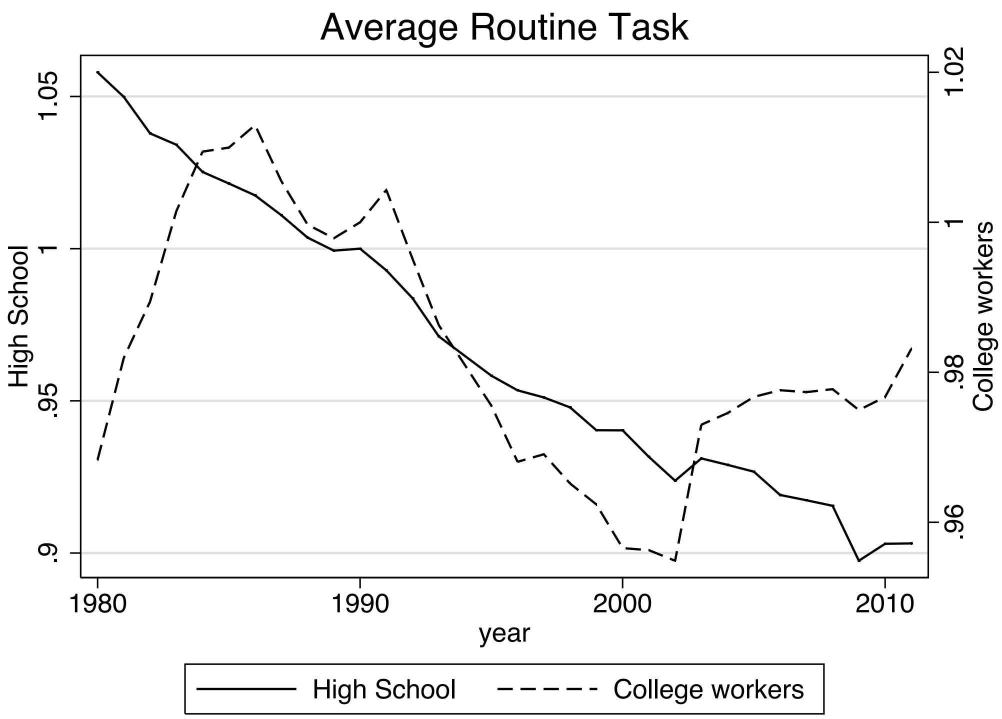

notes: The figure plots an index (normalized to one in 1990) of the average level of routine task for employed high school and college graduates over time using the CPS ORG data from 1980-2011. The routine task measure comes from the average of the variables `figure` and sts described in Autor, Levy, and Murnane (2003) and in the data appendix.

Figure 25:

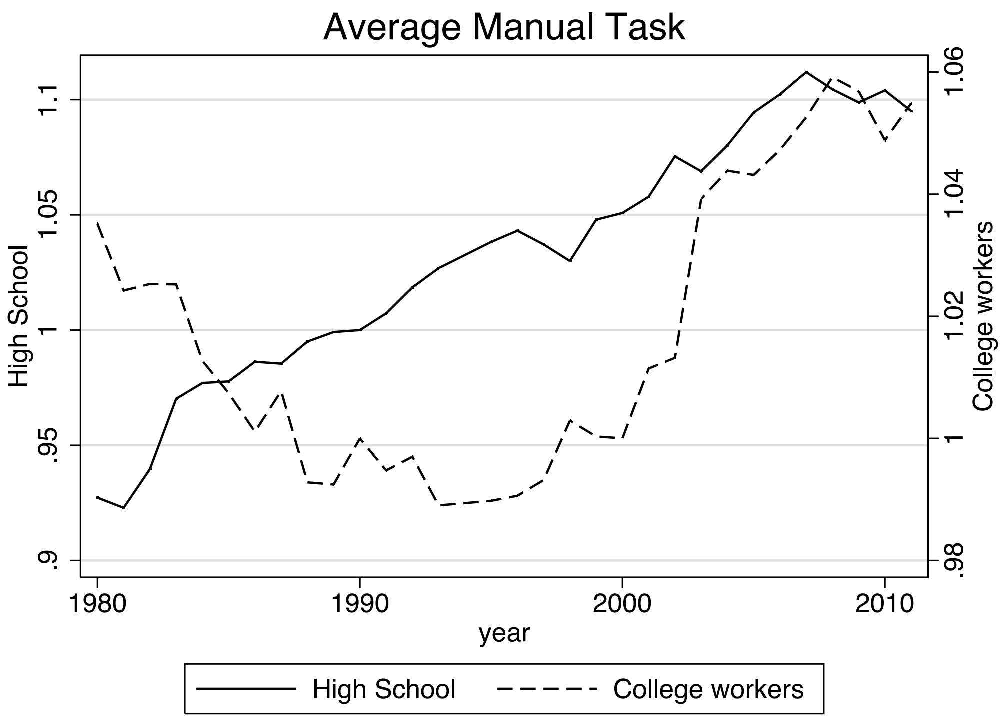

notes: The figure plots an index (normalized to one in 1990) of the average level of manual task for employed high school and college graduates over time using the CPS ORG data from 1980-2011. The manual task measure comes from the variable `ehf` described in Autor, Levy, and Murnane (2003) and in the data appendix.

Figure 26:

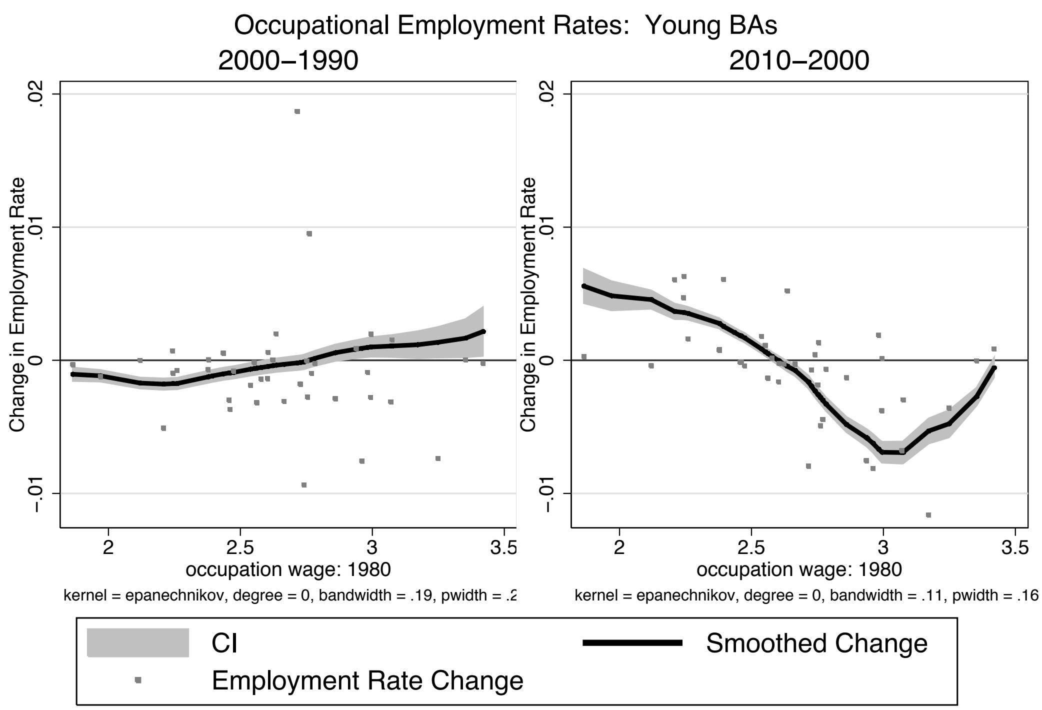

notes: The figure plots the change in employment rate for young college workers over the 1990-2000 and the 2000-2010 period by occupation against the occupation 1980 log wage. The solid line represents an estimated local-mean smooth of the employment rate changes.

Figure 27:

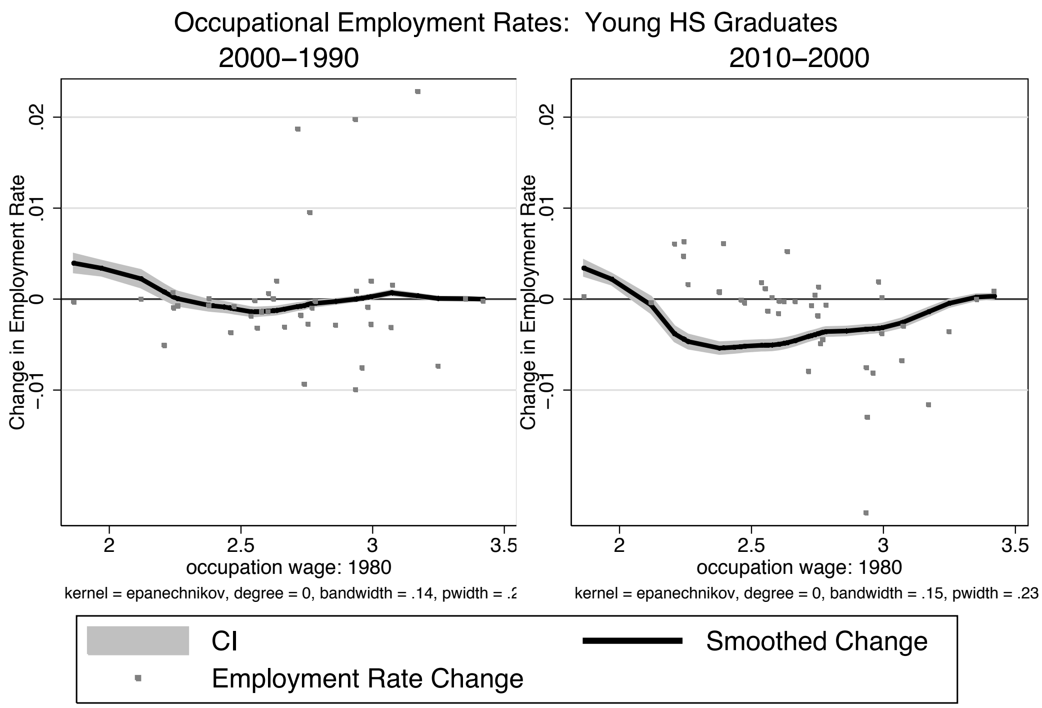

notes: The figure plots the change in employment rate for young high school graduates over the 1990-2000 and the 2000-2010 period by occupation against the occupation 1980 log wage. The solid line represents an estimated local-mean smooth of the employment rate changes.

Figure 28:

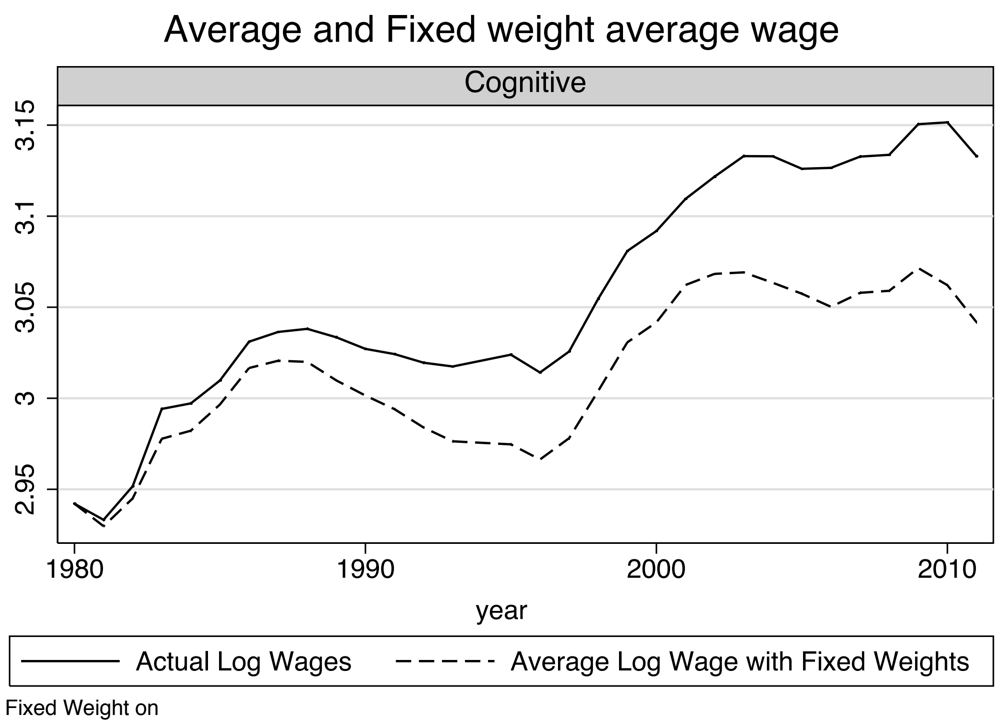

notes: The figure plots the average log wage by indicated occupation group by taking a raw average and an average using fixed weights, where the weights hold the demographic composition constant within an occupation over time.

Figure 29:

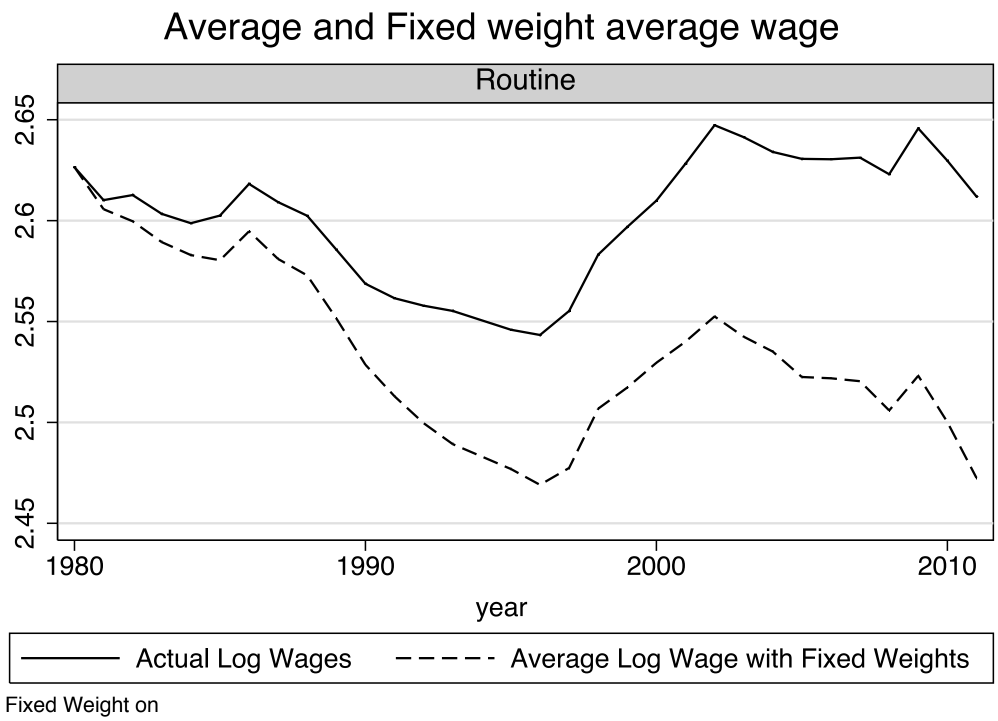

notes: The figure plots the average log wage by indicated occupation group by taking a raw average and an average using fixed weights, where the weights hold the demographic composition constant within an occupation over time.

Figure 30:

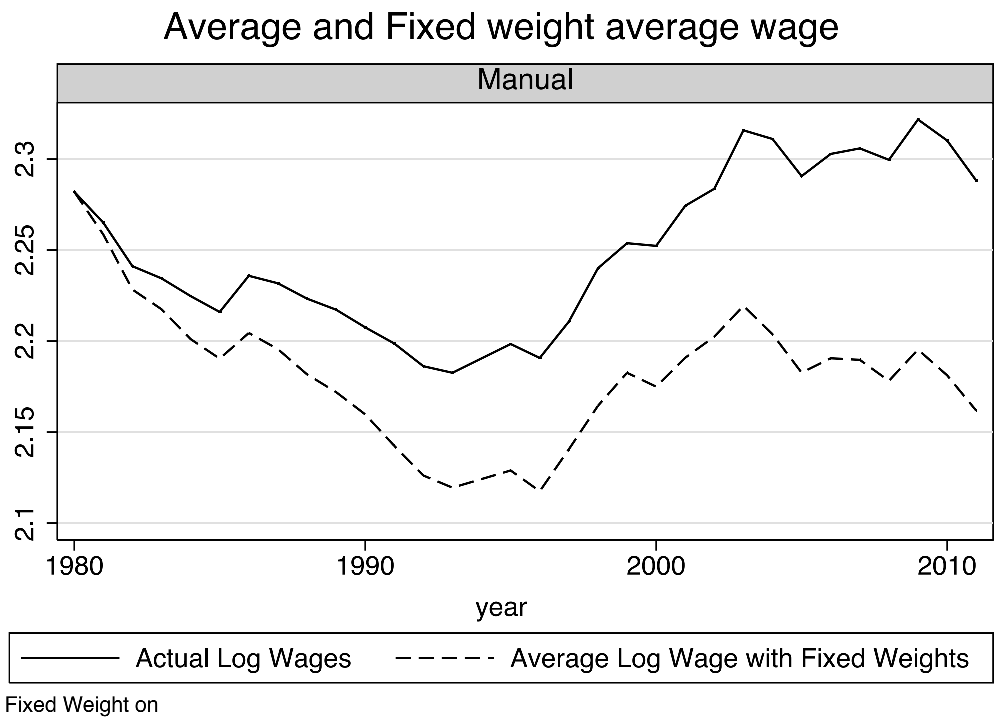

notes: The figure plots the average log wage by indicated occupation group by taking a raw average and an average using fixed weights, where the weights hold the demographic composition constant within an occupation over time.

Figure 31:

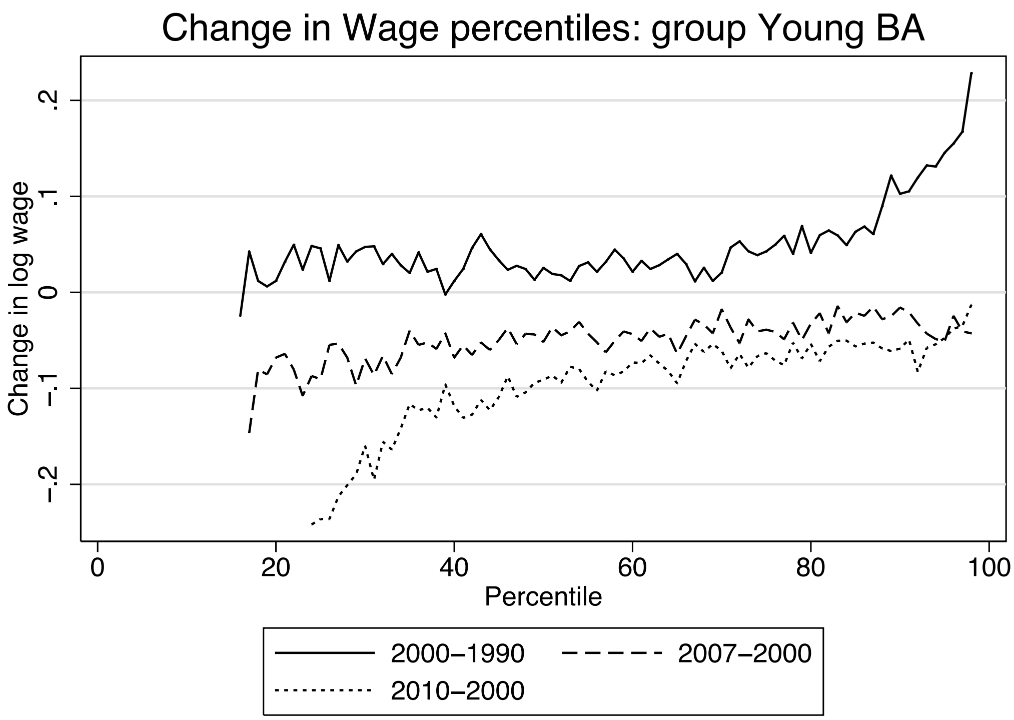

notes: The figure plots the change in log wage at each percentile in the wage distribution for three different time periods. When calculating the wage percentile, we include both allocated and non-allocated wages and do not remove outliers. For non-workers, we impute log wage as zero. The series are plotted starting at the first non-zero wage percentile.

Figure 32:

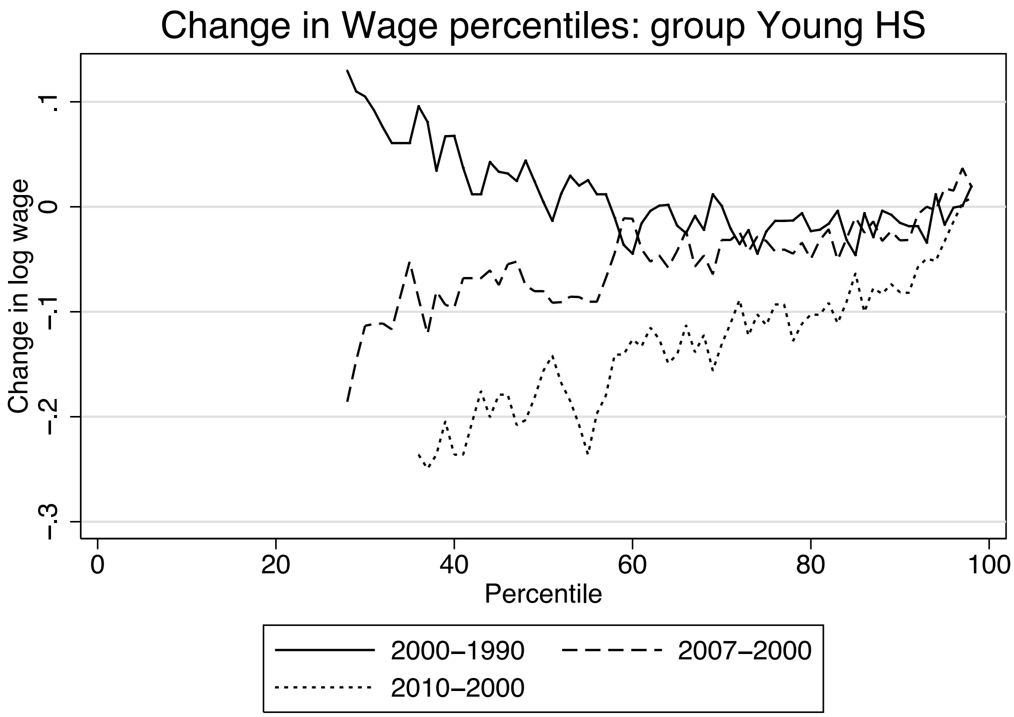

notes: The figure plots the change in log wage at each percentile in the wage distribution for three different time periods. When calculating the wage percentile, we include both allocated and non-allocated wages and do not remove outliers. For non-workers, we impute log wage as zero. The series are plotted starting at the first non-zero wage percentile.

Figure 33:

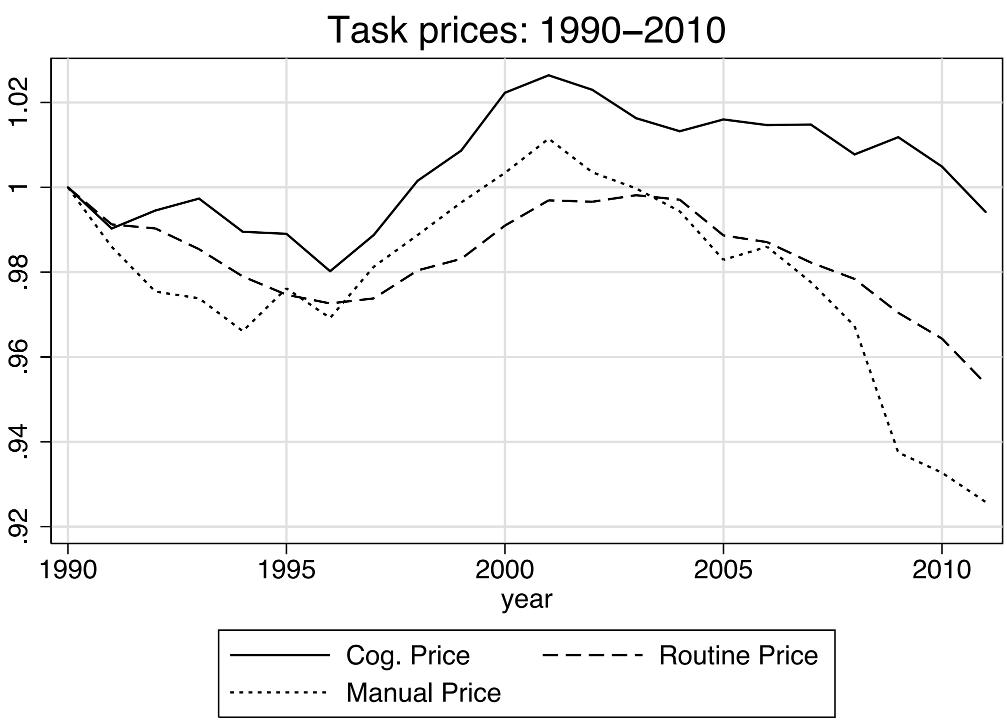

notes: The figure plots the task-price indices for manual, routine and cogntive tasks over time. The construction of the price indices is described in the text.

Figure 34:

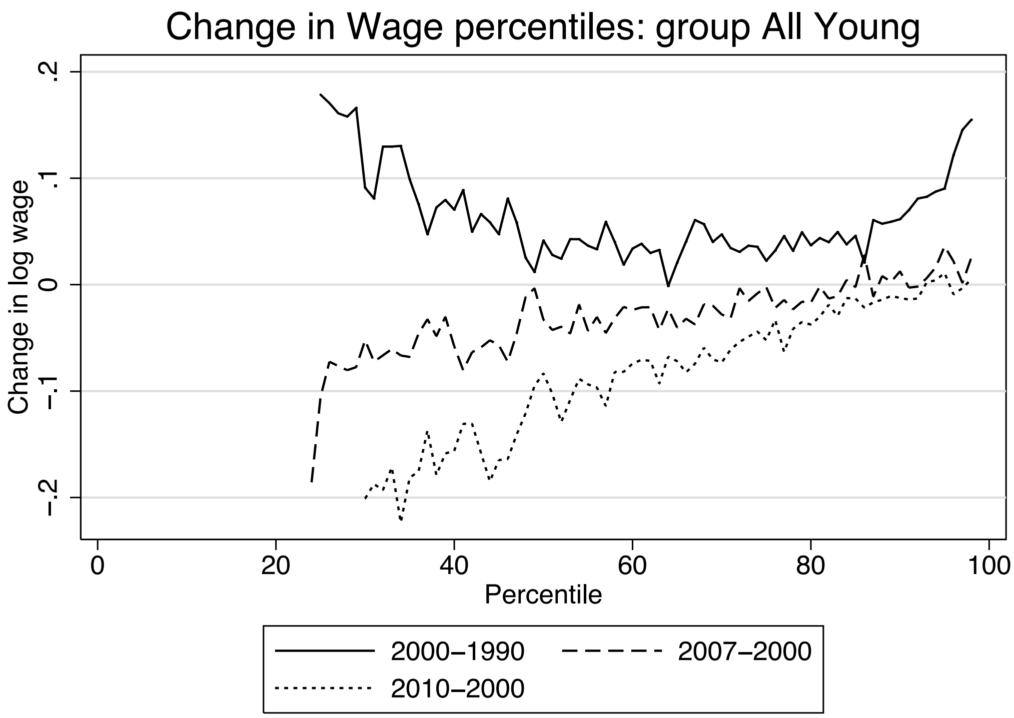

notes: The figure plots the change in log wage at each percentile in the wage distribution for three different time periods. When calculating the wage percentile, we include both allocated and non-allocated wages and do not remove outliers. For non-workers, we impute log wage as zero. The series are plotted starting at the first non-zero wage percentile.

Table 2: Estimates of Equation (8)

	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta_{2000-1990} M_c^{rate}$	-0.48*					
	(0.13)					
$\Delta_{2000-1980} M_c^{rate}$		-0.26*				
		(0.087)				
$\Delta_{2000-1990} \log M_c^{rate}$			-0.15*			
			(0.037)			
$\Delta_{2000-1980} \log M_c^{rate}$				-0.12*		
				(0.023)		
$\Delta_{2000-1990} ER_c$					-0.32*	
					(0.10)	
$\Delta_{2000-1980} ER_c$						-0.39*
						(0.055)
Constant	-0.017*	-0.011	-0.015*	0.0084	-0.034*	-0.016*
	(0.0053)	(0.0079)	(0.0054)	(0.0088)	(0.0020)	(0.0033)
Observations	231	231	231	231	231	231
R^2	0.071	0.060	0.089	0.13	0.069	0.24

Notes: U.S. Census and ACS data from 1980-2010. Unit of observation is the CMSA. All regressions are weighted by the square root of the city size in 1980. Robust standard errors in parentheses. (*) denotes significance at the 5% level.

A Data

A.1 May/ORG Current Population Survey

May CPS data from 1973-1978 and ORG CPS data from 1979-2011 are downloaded from the NBER[33]

- Initial extractions included all individuals between the ages of 16-64.
- Potential experience calculated as:

$$\max\left(\min\left(\text{age} - \text{years of school} - 6,\ \text{age} - 16, 0\right)\right)$$

- Sample further restricted to those with positive potential experience.
- Prior to 1992, education was reported as the number of completed years. In 1992 and after, education is reported in categories as the highest grade/degree completed.
  1. We convert categories to years of completed school in the post-1991 data based on Park (1994)
  2. We convert years into categories in the pre-1992 data based on Jaeger and Page (1996) (code provided by NBER).

A.1.1 Wage data

The construction of our wage data closely follows Lemieux (2006).

- Wage data is based on those who report employment in reference week.
- In all wage calculations, we set allocated wages to missing.
- Our hourly wage measure is based on reported hourly wage for those who report hourly payment and not adjusted for topcoding. For workers who are not paid hourly:
  1. We use edited weekly earnings. For the years 1984-1986, we use unedited earnings due to the higher topcode value.
  2. Adjust topcoded wages by a factor of 1.4.
  3. Divide the result by usual hours worked per week.
- For all wage data, we set to missing hourly wages below 1 and greater than 100 in 1979 dollars based on the CPI.[34]
- For all reported wage statistics, we construct a 'labor supply weight' by multiplying the usual weight in the May CPS and the earnings weight in the ORG CPS by usual hours divided by 35.
- We use these 'labor supply weights' when we construct occupational employment shares.

[33] Links are http://www.nber.org/data/cps_may.html and http://www.nber.org/data/morg.html
[34] CPI data from http://data.bls.gov/cgi-bin/surveymost?cu and includes all items.

A.2 Investment Data

The investment data we use was downloaded from the Federal Reserve Economic Data (http://research.stlouisfed.org/fred2/) on 2013-05-23. The series we use to construct our figures are:

1. Private fixed investment: Nonresidential: Information processing equipment and software (A679RC1Q027SBEA), Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate
2. Private fixed investment: Nonresidential: Information processing equipment and software: Computers and peripheral equipment (B935RC1Q027SBEA), Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate
3. Private fixed investment: Nonresidential: Information processing equipment and software: Software (B985RC1Q027SBEA), Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate
4. Private fixed investment: Nonresidential: Information processing equipment and software: Computers and peripheral equipment (B935RC1Q027SBEA), Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate
5. Gross Domestic Product, 1 Decimal (GDP), Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate

B Reweighting

Our reweighting procedure is based on DiNardo et. al. (1996). We chose a base year of 1989. We pool the base year with each year in our May/ORG data and construct a variable equal to one if an individual is observed in 1989. With this as the dependent variable, we run a logit regression. The right hand side variables include education (six categories), age (in two-year bins), and indicators for gender and non-white ethnicity. Education and gender are interacted with every variable.

We use the predicted values or propensity scores from these estimations to form counterfactual weights that hold the composition of the workforce constant over time. Our procedure closely follows Lemieux (2006).

C Occupation Categories

- The occupation categories we use are based on the 1980/90 Census occupation categories. Several small changes where made in the 1990 Census occupation classifications that required slight aggregation. We use the code provide by http://www.unionstats.com to make these adjustments.
- The categories are consistent from 1983-2002.
- For years prior to 1983 and after 2002, we use BLS crosswalks[35] to allocated workers to the 1980/90 categories.

[35] Obtained from http://usa.ipums.org/usa/volii/occ_ind.shtml

– Our broad occupation categories are made after converting all data to the 80/90 categories and aggregating up.

D Additional Figures: Temp section header for appendix tables

Figure 35:

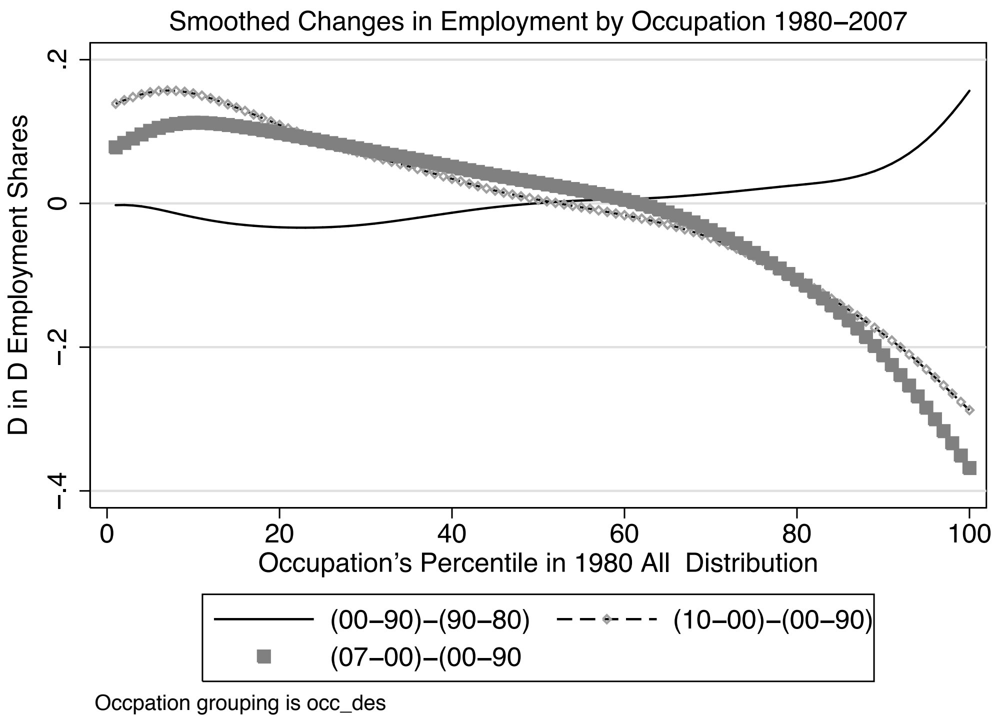

Figure 36:

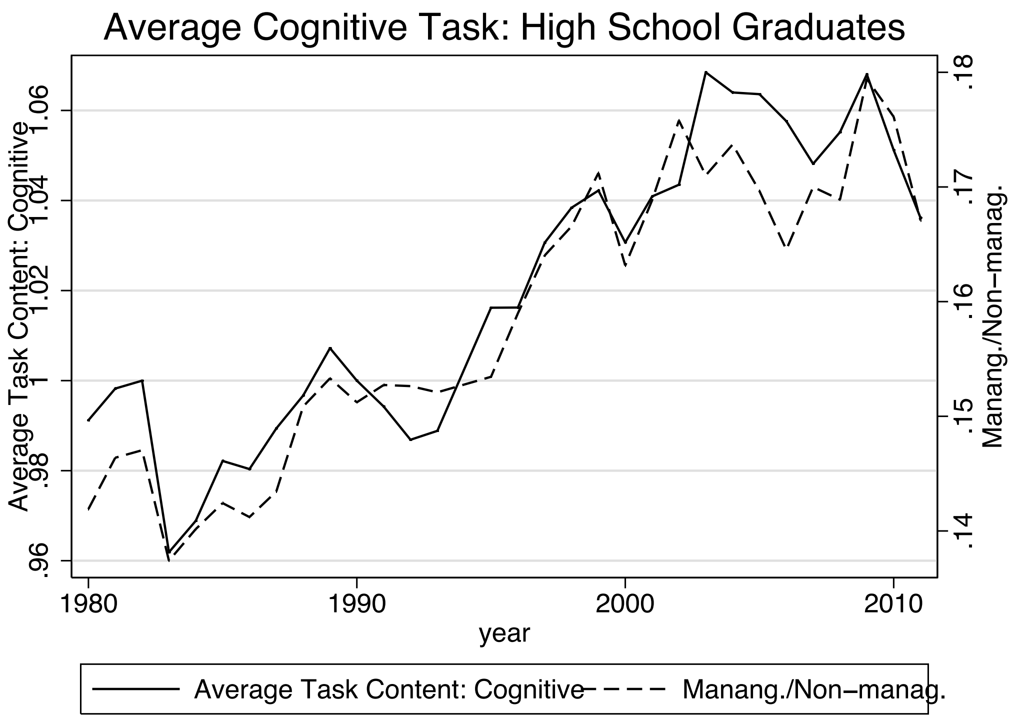

notes: The figure plots the average level of cognitive task and cognitive/non-cognitive employment ratio for high school graduates over time using the CPS ORG data from 1980-2011. The cognitive task measure comes from the average of the variables `math` and dcp described in Autor, Levy, and Murnane (2003) and in the data appendix.

Figure 37:

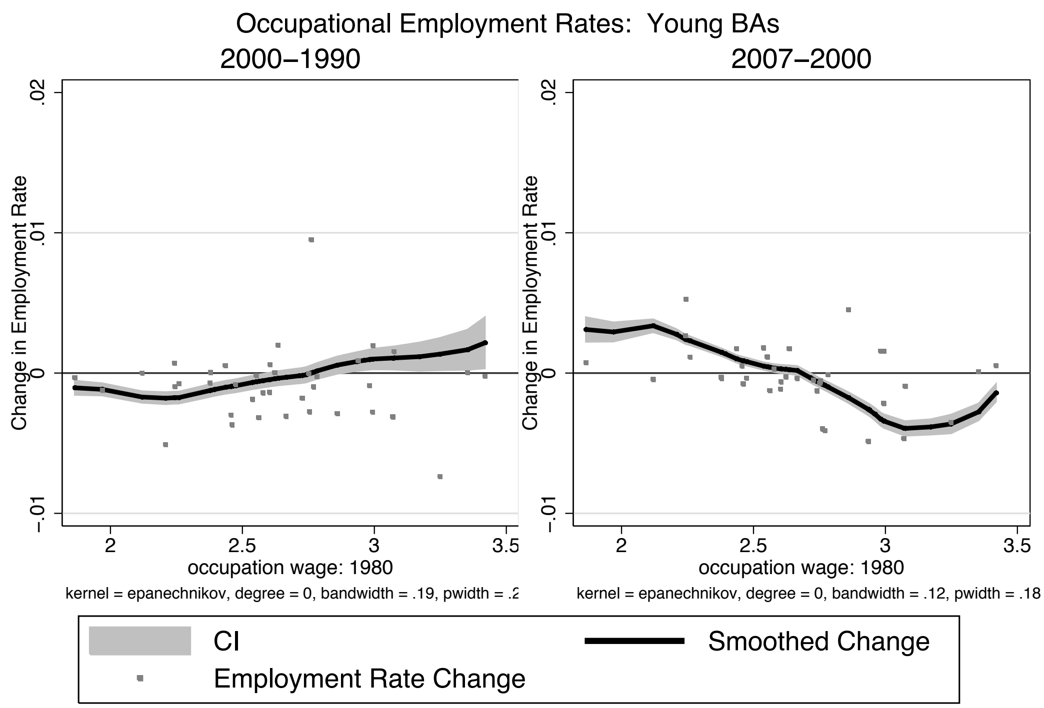

notes: The figure plots the average level of cognitive task for employed college graduates over time using the CPS ORG data from 1980-2011. The cognitive task measure comes from the average of the variables `math` and dcp described in Autor, Levy, and Murnane (2003) and in the data appendix.

Figure 38:

Occupational Employment Rates: Young HS Graduates
2000−1990
2007−2000
Change in Employment Rate
−.01
0
.01
.02
2
2.5
3
3.5
occupation wage: 1980
kernel = epanechnikov, degree = 0, bandwidth = .14, pwidth = .2
kernel = epanechnikov, degree = 0, bandwidth = .13, pwidth = .19
CI
Smoothed Change
Employment Rate Change

Figure 39:

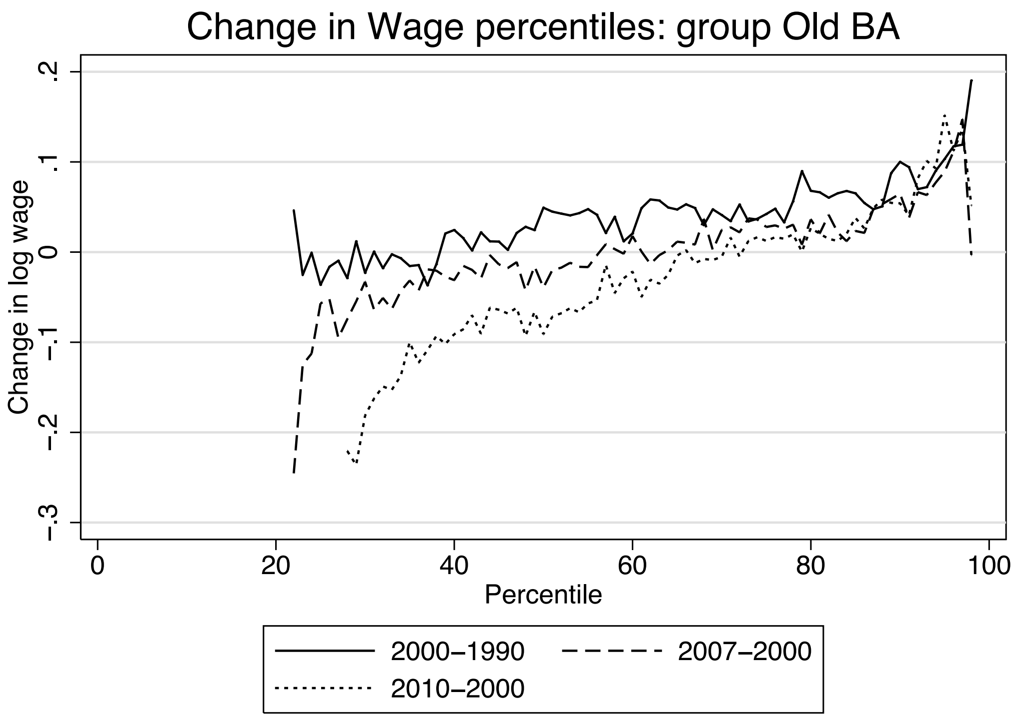

notes: The figure plots the change in log wage at each percentile in the wage distribution for three different time periods. When calculating the wage percentile, we include both allocated and non-allocated wages and do not remove outliers. For non-workers, we impute log wage as zero. The series are plotted starting at the first non-zero wage percentile.