The Great Reversal in the Demand for Skill and Cognitive Tasks

Paul Beaudry

University of British Columbia and NBER

997-1873 East Mall, Vancouver, BC Canada V6T 1Z1.

Paul.Beaudry@ubc.ca

David A. Green

University of British Columbia and Research Fellow, IFS, London

997-1873 East Mall, Vancouver, BC Canada V6T 1Z1.

David.Green@ubc.ca

and

Benjamin M. Sand

York University

1090 Vari Hall, 4700 Keele St., Toronto, ON, M3J 1P3.

bmsand@yorku.ca

This paper argues that several of the poor labor market outcomes observed in the great recession can be traced back to a change in the demand pattern for skilled workers that started with the tech bust of 2000. In particular, we show that around the year 2000, the demand for cognitive tasks underwent a reversal. In response, high-skilled workers moved down the occupational ladder and increasingly displaced lower-educated workers in less skill intensive jobs. While these effects where present before the financial crisis of 2008, they became more obvious after jobs associated with the housing bubble disappeared.

The poor performance of the US labor market since the financial crisis of 2008 continues to generate enormous debate around 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 de-leveraging 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 et al. (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 et al. (2003); Autor and Dorn (2013); Autor et al. (2006, 2008); Dustmann et al. (2009); Firpo et al. (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 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. Such a decline in demand has had, and continues to have, a direct impact on more skilled workers, but 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, 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. 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 along with it, can be seen as a General Purpose Technology (GPT) and, like all GPTs before

[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.

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 – it is also about beating out less educated workers for barista and clerical-type jobs.

At first pass, the idea that the IT GPT reached a sort of maturity stage around 2000 may seem surprising. The theory does fit with overall wage and employment patterns, but we also corroborate it with two other pieces of evidence. First, we will show that investment in both IT equipment and software showed sharp breaks in 2000, moving from previous periods of strong growth to either declining or stagnant trends after 2000. Thus, there is reason to believe that something changed with respect to IT implementation at the time of the tech-bubble crash and that the change has persisted. Second, we investigate patterns of change within and across industrial sectors. In that exercise, we find that the majority of the change in the cognitive task employment pattern between the 1990s and 2000s happened within all industries (which fits with the idea of IT being a GPT). At the same time, for college educated workers, the management services industry stands out as a sector that underwent very substantial growth in employment in the 1990s and then decline in the 2000s. We show that if this sector alone had continued its 1990s growth pattern in the 2000s, then the post-2000 decline in the proportion of young college workers employed in cognitive occupations would not have happened. Since this sector includes both management services and computer related services, its importance and pattern fits with idea of IT being a GPT that went through an investment stage before 2000 and a maturity stage thereafter. Many industries working to incorporate this new technology would have turned to consultants while putting the new technology in place and those consultants would show up in the management services sector.

Our second main contribution is that we provide a detailed picture of the changes in employment and wage patterns 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. Section 4 concludes.

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 1979-2013. 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 when calculating wage statistics.[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 our main empirical work, we restrict the sample to

[2]The Data appendix provides additional information on our data processing.

individuals aged 25-54 with positive potential work experience (i.e., with age - years of schooling - $5 > 0$) to ensure that our patterns are not being driven by schooling decisions at the young end or changes in retirement at the older end. In several instances, we will present results just for young individuals, by which we mean individuals aged 25 to 35. We will also present some results broken down into two education groups: College graduates, which includes all individuals with a college or higher university degree; and high school graduates.

1.2. Wage and Employment Rate Patterns

In Figure 1 we plot the employment rate of individuals aged 25 to 54. 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 substantially over the 1980s and 1990s, apart from the downturn in the early 90s recession, 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. We interpret the graph as showing a trend break in 2000, and our goal is to explore the nature and potential causes of that trend break. One could, alternatively, view the graph as showing a relatively stationary series with a larger than normal boom peaking in 2000 and bracketed by two historically large recessions. Under our interpretation, the prognosis for the immediate future is pessimistic, while under the second interpretation, we are just witnessing cycles of varying amplitudes and one would predict a strong period of growth in the near to medium-run future. Later in the paper, we will present investment series that we see as fitting with the trend break interpretation but, of course, we have no way of being sure which forecast for the near future is the right one.

In Figure 2 we plot the average real hourly wage for all workers against the employment rate. We plot the post 1990 period to focus attention away (to some extent) from the role of the increased labor force attachment of women. Based on this figure, it is not the employment rate alone that changes after 2000, but the whole relationship between employment and wages. From the end of the 1990s recession to 2000, both employment rates and real wages rose substantially. But after 2000, the relationship changes direction, with employment rates falling while real wages climbed. Even with the (somewhat small) improvements in the employment rate after 2011, real wages remain over 7% higher than their levels in the early 1990s. Interestingly, the 2003-7 period appears as what might be characterized as a stalling phase in the general pattern. This fits with the analysis in Charles et al. (2012), who argue that the construction boom served to mask longer run negative employment patterns for non-college men. In a simple demand and supply framework, the post 2000 pattern would fit with an inward shift of the labor supply curve. While some authors have suggested such a shift in response to enhanced UI benefits after 2008, we do not know of

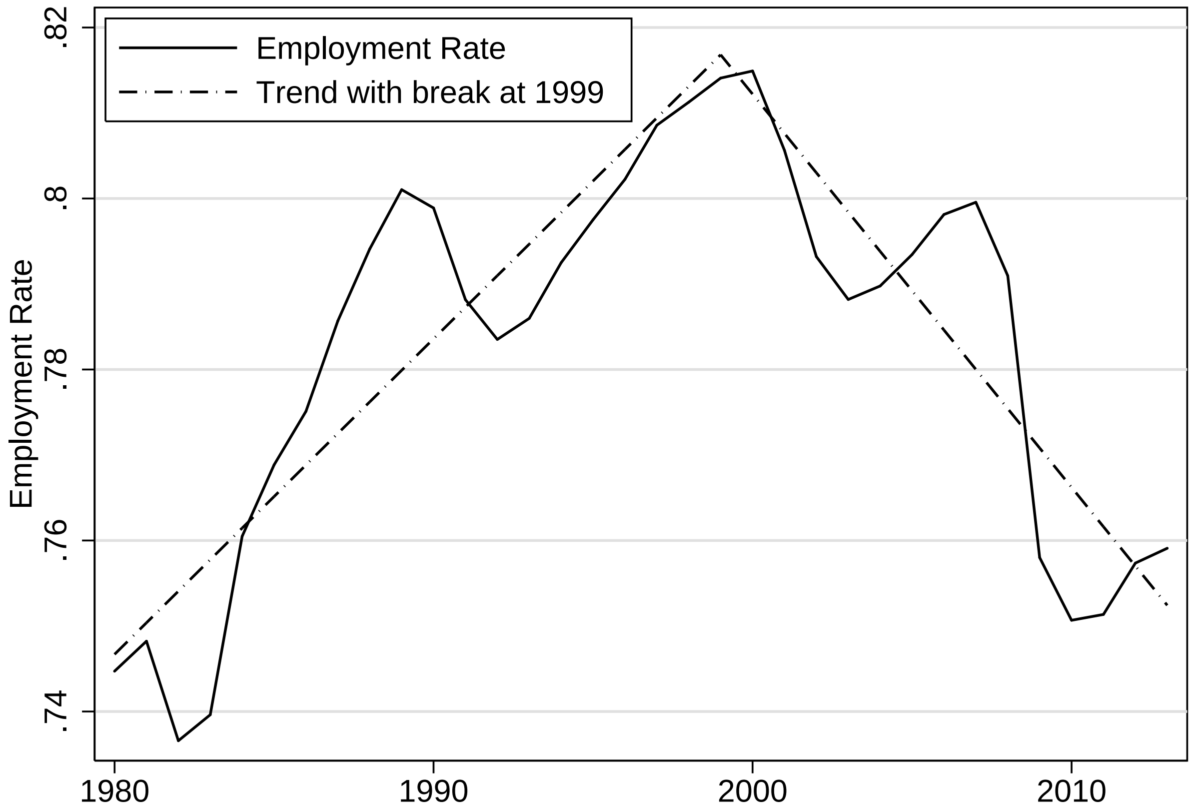

Fig. 1.— Employment Rate with Fitted Trend

a candidate explanation for why such a shift would have occurred starting in 2000. In the next section, we present a theoretical model which can explain the post-2000 pattern without a supply shift.

Plotting the same figure for men and women, separately, and for different education groups (not shown here), we find the same general pattern with the same directional change in the wage-employment relationship in about 2000 in each case. The main exception is for the high school educated, where the employment rate follows the same pattern seen in Figure 2 but real wages fall back to their early-90s level by 2013. The fact that the high school employment rate is 7 points lower in 2013 than in 1996 (0.71 versus 0.78), even after real wages are back at the same level suggests that there has been a structural change of some kind in the intervening years. At the same time, the fact that the same general pattern is observed for men and women, and for the highly educated suggests that the pattern in 2 is not being driven by easily observable shifts in composition.

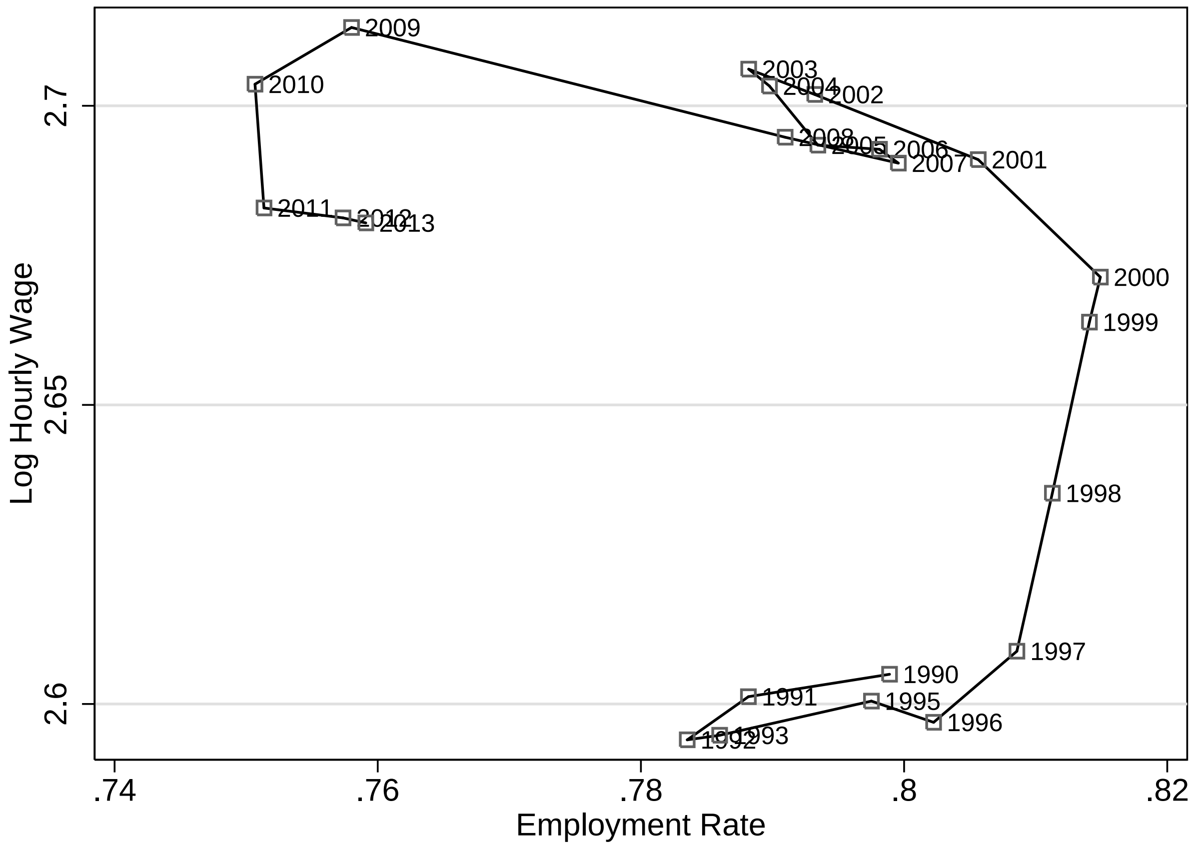

Fig. 2.— Log Hourly Wage vs Employmen Rate

While both the college graduates and high school educated groups show the same general employment rate patterns, their movements are far from identical. In Figure 3 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 increased faster than that of the college workers, while after 2000 we see 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 three figures motivate us to downplay the housing boom and bust of 2003-2008 for understanding the current labor market situation and to focus, instead, on longer-run trends. In particular, we want to ask: “What type of force could be driving these medium run patterns in which employment and wages increase for at least a decade then switch direction with employment falling but real wages staying high or rising?” We are interested, moreover, in whether there is a force that could explain both the pre- and post-2000 patterns in a unified way. It would always, of course, be possible to explain the patterns with different forces operating before and

Fig. 3.— High School – College Employment Rate Ratio

after 2000 but such an approach would impose little discipline on our investigation and would still call for an overall explanation for why the driving force changed in 2000.

Since 2000 was the year of the bust in the high-tech sector, forces related to the IT revolution seem to have the potential to provide a unifying explanation. To investigate whether that might be the case, we turn to plots of IT related investment in Figure 4. The panel (a) in the figure corresponds to investment in information processing equipment and software as a ratio to GDP for the US. With the exception of a slowdown in the late 1980s, the series shows a prolonged and substantial increase between 1960 and 2000. After 2000, however, this broad IT investment rate first jumps down then follows a declining path. The panel (b) focuses on the hardware part of this pattern, containing a plot of the ratio of investment in computers to GDP. This series also shows the increase to 2000 and the jump down immediately after the tech bust, with a more dramatic decline in the ensuing years. Finally, in panel (c), we plot the ratio of investment in software to GDP. In this case, there is no evidence of a post-2000 decline but there is a clear trend break in 2000, with the series being quite flat in the 2000s. We view these series as evidence that forces related to the IT revolution are a reasonable place to focus attention in looking for an explanation for the 2000 break in the wage and employment series. In the next section, we present a model based on that idea and turn, in the ensuing sections, to investigating the empirical implications of that model.

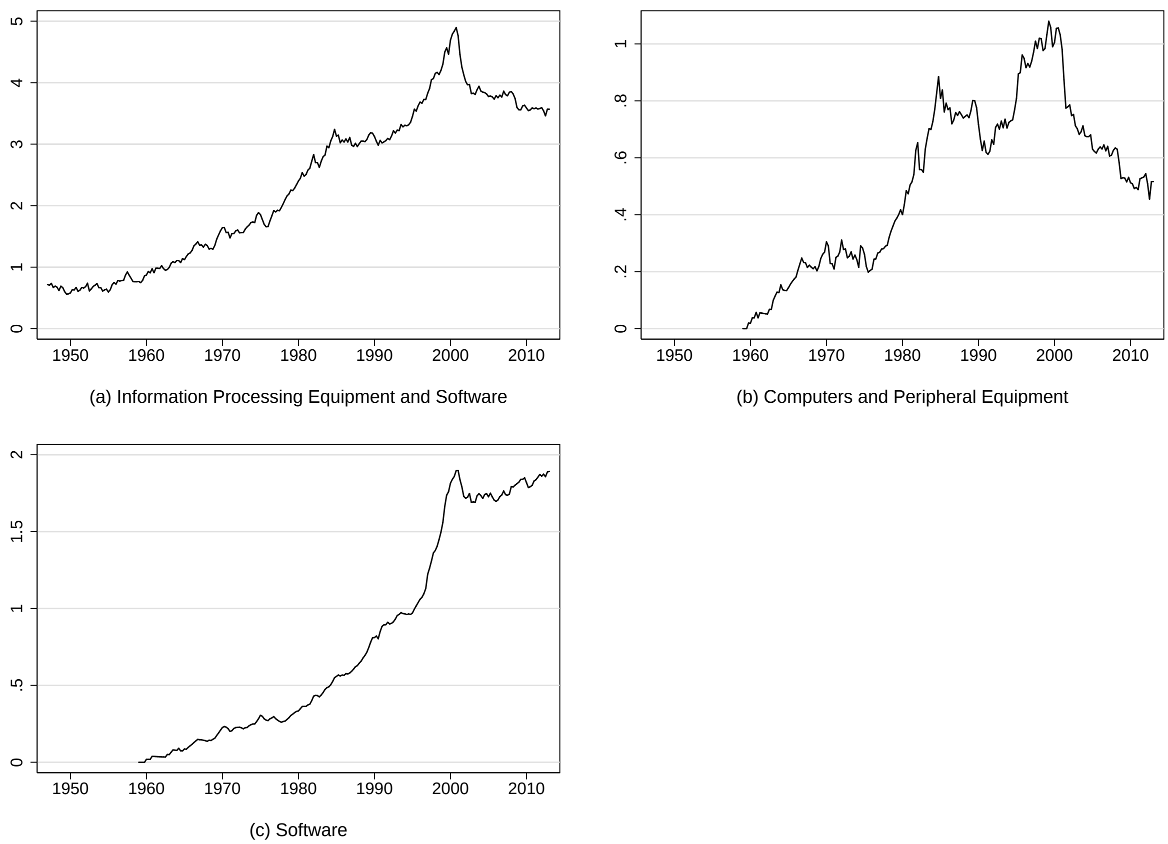

(a) Information Processing Equipment and Software

(b) Computers and Peripheral Equipment

(c) Software

Fig. 4.— Investment in IT as a share of GDP

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. While we will present some comparative static results, our main interest will be in the transitional dynamics that arise after the initiation of a technological change. In this sense, our approach is very much in the spirit of papers such as Welch (1970), Schultz (1975) and

Bartel and Lichtenberg (1987) which emphasize the notion that skills are particularly important during adjustment phases. In those papers, more educated entrepreneurs (where Schultz (1975) defines this term quite broadly) imply more efficient uses of other inputs to production and more effective adoption of new technologies. Similarly, in our model, cognitive skills are needed to get a new technology working.

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 them 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, $F_{L^r,\theta} = 0$.[3] 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 starting at least as far back as Autor et al. (2003).

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 or 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-educated 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

[3] 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.

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.[4]

The decision problem for less educated workers is very similar. [5] 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. 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. As with the highly educated, the decision problem of the less educated workers is characterized by a cut-off $\bar{\phi}_t$, where all workers with $\phi \leq \bar{\phi}_t$ supply their labor to cognitive tasks. For both types of workers, the proportion choosing the cognitive sector will increase if w_t^c increases relative to w_t^r.

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 (i.e., the employment rate) will be determined solely by demand. Because 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 the second ($w_t^r = A$) equilibrium holds. In this case, the employment of routine tasks is implicitly determined by $F_{L^r}(\Omega_t, L_t^r, N, \theta_t) = A$. We assume that routine employment for less educated workers is this demand minus the amount of the routine task supplied by more educated workers.[7]

---

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

[5]In the appendix we provide complete details on the model and its solution.

[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 $b \cdot (1+\alpha) < 1$.

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:

– 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.[8]

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

[8] 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.

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.

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. In the appendix, we provide theoretical results based on specific functional forms. This permits analytic solutions which clarify insights into the main model implications.

The dynamics in the model involve an initial stage where the 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.[9] 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.[10]

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.

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

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

– 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. It is this technology implementation stage that is most reminiscent of the models in Welch (1970) and Schultz (1975).

However, as the next proposition indicates, this boom will eventually be 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.[11]

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.

– Even during the 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,

[11] 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.

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 continues to be replaced by the improved organizational capital. One group continues to do well throughout the two phases, though: the entrepreneurs. They are the ones reaping the benefit of the technological advance, fitting in a broad way with the pattern of larger shares of income going to the top 1%.

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.[12] 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 de-couple movements in the skill price from movements in the observed average wage in cognitive occupations. When the price of effective units goes up, 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

---

[12] 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 $2A$.

whole economy behaves during the bust period.[13] 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 5 plots the joint movement of the average wage and the employment rate for a simple parametrization of the model.[14] As can 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, with the boom period corresponding to the 1990s and the bust period being the 2000s. The model implies that the muted movements in the overall average wage in the 2000s are a reflection of selection. 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 Figure 2. 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. Introducing 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 particular, it is overly simplified in including only two types of tasks

---

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

[14] 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 $\psi \in [0,1]$ 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.

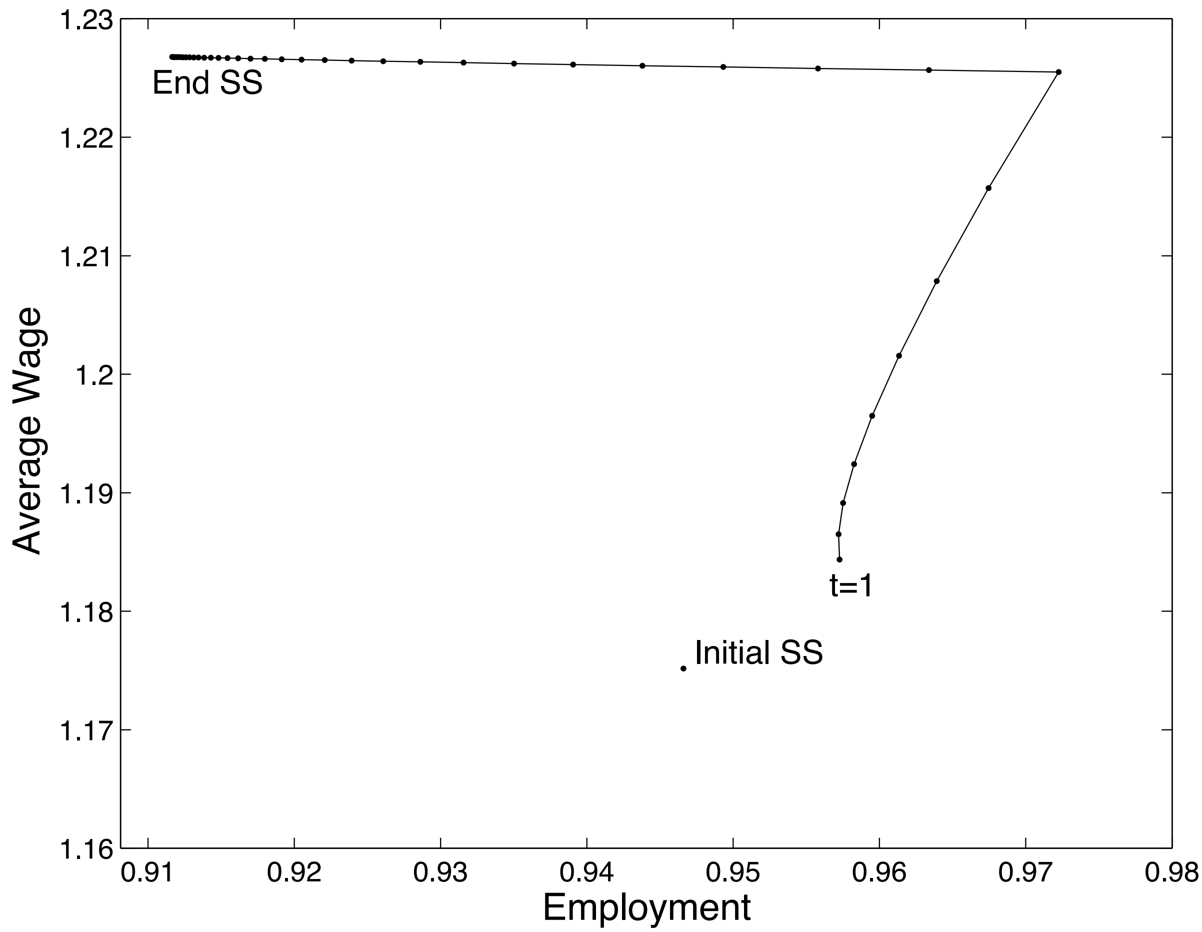

Fig. 5.— Employment Dynamics from a Simulated Model

(Cognitive and Routine). The characterization of production in what might be seen as the goods producing sector as just using these two types of tasks is common in the polarization literature (e.g., Autor and Dorn (2013)). The lower paid service sector (which employs non-routine manual tasks) is then added on by, for example, having utility being a function of both goods and services. We left the service sector out of our model 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.

Including a manual, non-routine (i.e., service) sector in our discussion leads to the following heuristic description of the impact of the technological revolution. In the boom period (the 1990s), there is shift out in the demand curve for cognitive tasks. In our model, this happens because firms are building organizational capital to take advantage of the technical change. The resulting rise in the cognitive task price implies a move up the supply curve for these tasks. At the same time, in the market for routine tasks there is 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.[15]

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 (i.e., the demand curve shifts out) because of increased demand for services coming from segments of the population that are becoming richer. 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. If the net effect is small, the increased demand for manual work implies the price of these tasks will rise. This, in turn, would favor the movement of people from the household sector to the manual task sector. 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.

In the bust phase (the 2000s), the demand curve for cognitive tasks shifts down relative to the 1990s as the demand for cognitive workers to implement the new technology dissipates. In the routine 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. 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. The net result 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 supply: more employment but lower skill prices. This contrasts with the evidence of a positive shift in demand

---

[15] 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.

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; 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 1980-2013 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 6 we plot the employment rate of individuals aged 25-54 employed in occupations that require substantial cognitive skills for each year from 1980 through 2013. The employment rate in the figure is calculated as the total hours worked in cognitive jobs over the size of the workforce and represented as the log-change from 1980. This ratio increased substantially between 1980 and 2000, as one would expect given all the attention that has been paid to skill biased demand shifts. But, remarkably, the ratio flattens after 2000. In the same figure, we also plot the growth what could loosely be called 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, and plot the percentage growth in this index from 1980. It rises strongly over time because of increased education levels in the labor force (among other reasons). Essentially, it says that given the rise in education levels and given the occupations that more educated people tended to work in during the base year (1980), one would predict an increasing proportion of the labor force would see itself as suited to cognitive task jobs.

There are two key features of Figure 6 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. This constitutes a first piece of evidence suggesting an important reversal in the demand for cognitive tasks beginning in the year 2000.

It is interesting to consider the compositional underpinnings of the post-2000 slowdown in cognitive occupational growth. To that end, in Figure 7, we plot changes in the cognitive employment rate from 1980 for high school graduates (panel (a)) and those with a college or higher degree

Fig. 6.— Cognitive Employment Rate and Supply Index

(panel (b)), separately, and show the movements in occupational sub-groups. Overall, employment in cognitive task jobs relative to the population for high school graduates continued to increase until 2004, trending sharply downward after that point. For the university educated, the decrease occurred starting in 2000. For the high school graduates, the rise and fall are larger (notice the y-axes labels) and are almost entirely accounted for by employment in management occupations. For the university educated, management also plays an important role in the overall pattern but the Math, Science and Engineering category also exhibits a pattern of growth to 2000 followed by decline or stagnation.

In Panel (a) of Figure 8, we plot the log-change in the employment rate and supply index for routine (i.e., production and clerical) occupations. The employment pattern follows the observation in Siu and Jaimovich (2012) that the decline in routine task employment since the 1980s in the US occurred in a ratcheting pattern with declines in recessions that are not recouped in subsequent booms. The supply index indicates that as education levels rose, the proportion of people who we would expect to occupy routine task jobs declined. This decline essentially matched the fall

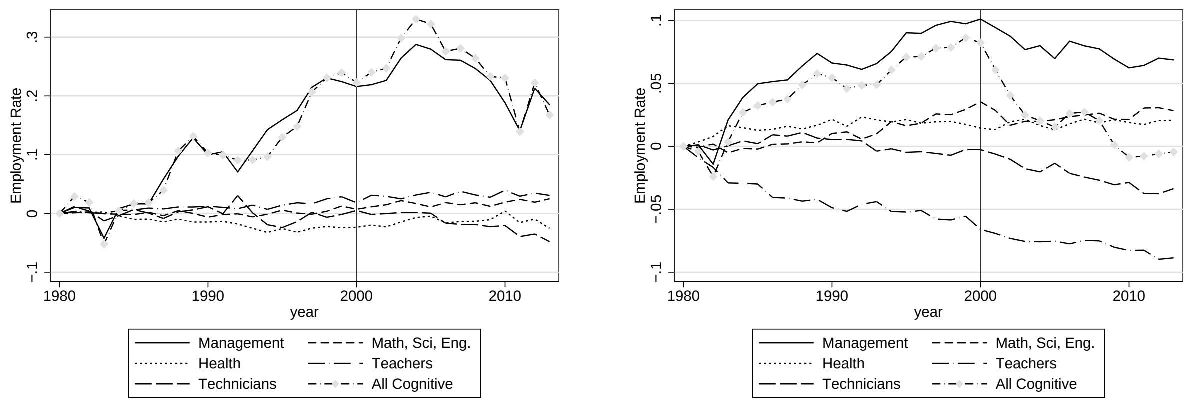

(a) High School Graduates

(b) College Graduates

Fig. 7.— Employment Rates for Cognitive Occupations

in employment up to 2000, suggesting that demand and supply in these occupations were falling at a similar pace. But after 2000, employment fell dramatically. Our model does not necessarily predict this sharp drop since it implies a combination of a drop in demand and an increase in supply as workers exit the cognitive sector. Instead, the timing may fit with a combination of technical change forces and the impact of increased Chinese imports emphasized in Autor et al. (2013).

(a) Routine Occupations

(b) Manual Occupations

Fig. 8.— Employment Rates and Supply Index: Routine and Manual Occupations

Finally, in Panel (b) of Figure 8, we consider manual jobs. Employment in these occupations grew steadily, if slowly, throughout the period, with a surge upward in the early 2000s. In contrast, the proportion of people one would predict to be in these jobs declined strongly over time as education levels rose. Our model implies that the increase in employment in these service occupations relative to the supply of less educated workers reflects a demand shift in favor of these occupations before 2000 but a supply shift as more educated workers moved into these occupations after 2000. If this is true, we would expect to see an increase in service sector wages before 2000 but a decline after 2000. We return to this point in our consideration of wages, later.

One drawback of the approach we have taken so far is that it involves a somewhat artificial division of occupations into a small number of groups. But occupations typically 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 et al. (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 log-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.[16] Each year of data in the figure actually represents two adjacent years that we pool together to reduce noise. Figure 9 contains the resulting smoothed employment growth changes for the periods 1980-1990, 1990-2000, 2000-2007, 2000-2010.[17]

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 et al. (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 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. In the figures, 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. A comparison of these lines indicates that 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. Thus, we see a dominant feature of this figure being the evaporation of the growth in employment in top paying jobs after 2000. We see this as implying a'great reversal' in the demand for cognitive tasks. Later, we argue that the wage patterns are supportive of the notion

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

[17]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.

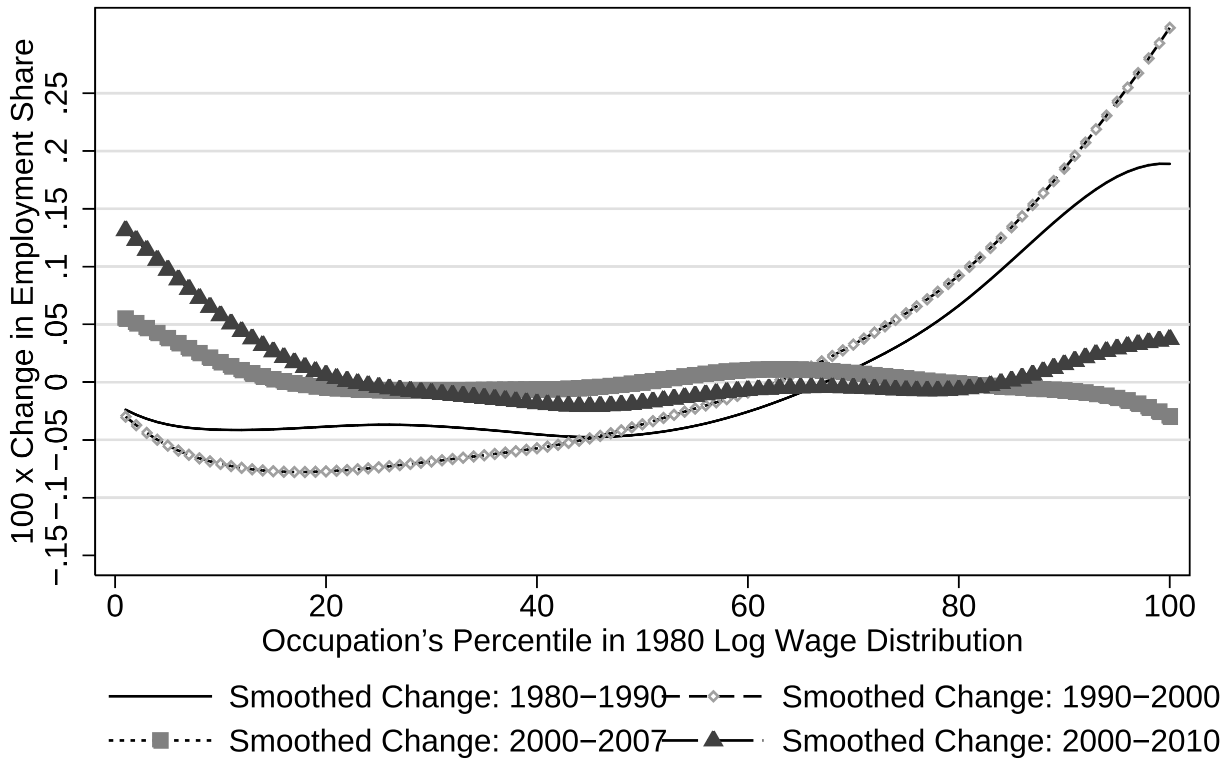

Fig. 9.— Smoothed Changes in Employment by Occupation: 1980-2010

of such a reversal.

3.2. Skill Upgrading and Downgrading

In the previous section, we presented evidence suggesting that the demand for cognitive task skills reversed after 2000. Our model fits with this pattern but also includes the prediction that workers in all education groups will move into cognitive task occupations and out of them in the 2000s. Moreover, we hypothesized that this will follow a "cascading" pattern in which the most educated will move into lower paying routine and manual jobs after 2000, pushing the less educated to even lower paying jobs and/or out of the labour market. We already saw evidence of the in and out nature of the flows to cognitive occupations for both high school and university educated workers in Figure 7. We pursue this further, and show patterns for the other occupation groups, using an alternate measure in Figure 10. Panel (a) in this figure contains plots of indexes of average cognitive task intensity of employment for high school and university graduates.

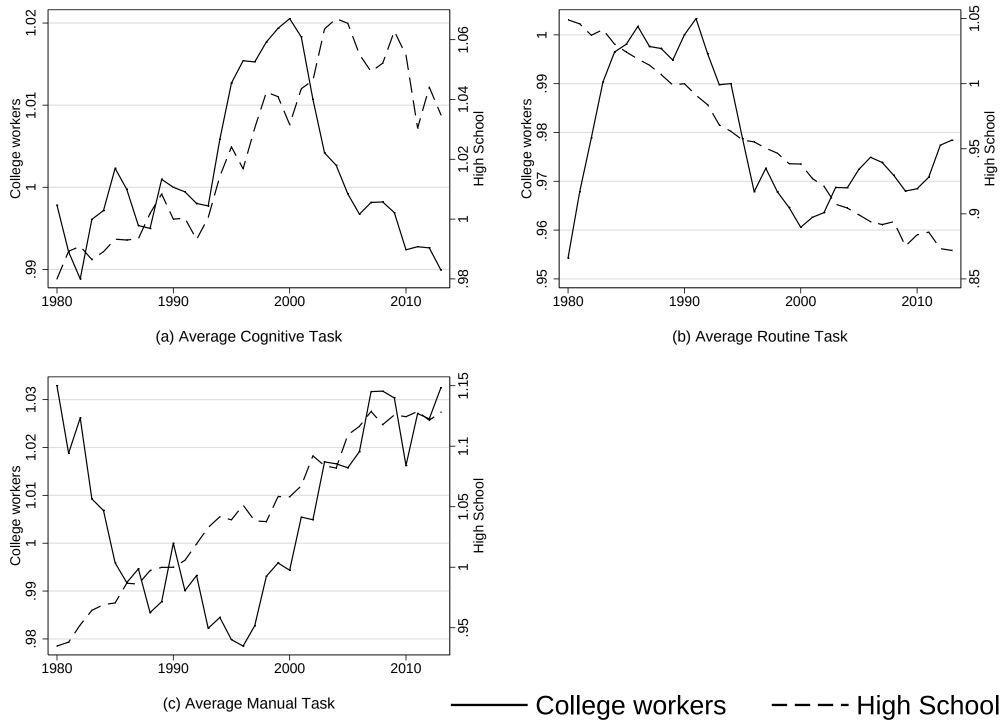

Fig. 10.— Occupation Task Content. Left scale refers to college graduates and right scale to high school graudates.

To construct the index, we first assign to each 3-digit occupation the average of the scores for cognitive tasks for that occupation from the Dictionary of Occupation Titles (DOT). We define cognitive tasks as the non-routine analytic and interactive tasks used in Autor et al. (2003) in their examination of the skill content of jobs. We then get the overall index value for a year by taking the weighted average of these occupation specific scores, where the weights are the proportion of employed people in each occupation. Thus, movements in the index reflect movements in college educated workers across occupations. Since this index takes account of the cognitive content of all possible occupations, it provides a more general measure of engagement with cognitive tasks than the approach in the last section of focussing on a subset of occupations. We normalize the index to equal one in 1990 so that movements in it can be interpreted as changes in average cognitive task usage relative to that year. The remaining panels contain plots of indexes constructed in the same way for routine and manual tasks.[18]

[18]We use data on the content of five different tasks for occupations coded at the 1980/90 Census Occupation Code level from Autor et al. (2003). We use these data to calculate three task measures that are closely related to measured used in Autor et al. (2003). First, we calculate a Cognitive Task measure by using the DOT codes `GED Math`, which measures general education development in mathematics and measured on a scale where a higher number indicates

For college educated workers, this set of graphs indicates a movement into occupations with greater cognitive task content in the 1990s and a move out in the 2000s. This matches what we saw earlier in terms of the subset of high cognitive task content jobs. Offsetting this was a movement out of routine (largely clerical) and manual task jobs before 2000. After 2000, college educated workers moved back into both routine and manual intensive occupations. The routine movement is particularly interesting since it is the opposite of the long term trend out of these types of occupations that can be seen in the high school grads data. For the latter group, the dominant pattern is a movement out of routine production and clerical occupations and into service/manual occupations. They also move into and then out of cognitive task occupations, but the latter are never represent a large proportion of employment for this group.

To obtain a more detailed view of the occupational shifting by different education groups, in Figure 11 we plot smoothed decadal changes in employment shares for occupations ranked by their average wage in 1990, in a manner similar to Autor et al. (2006)'s investigation of changes in the occupation structure. We do this separately for high school and college graduates. Occupational shares are constructed by calculating the share of hours worked in an occupation for a given group. Taking the difference in occupational shares over time, when the occupations are ranked by their average wage in 1990, indicates whether a given education group has shifted employment systematically in terms of occupations ranked by wage. For the high school graduates, the 1980s reflects a loss of middle-paying occupation jobs being offset by an increasing number of low-paying occupations jobs. In the 1990s, this pattern is muted, with less job loss in the middle, less gain at the bottom end, and some increase in employment in top-paying occupations. From 2000 to 2010, however, there is a loss in employment over a broad range of middle- and high-paying occupations with an offsetting growth in the bottom three deciles. This fits broadly with a picture of improvements in the middle and top in the 1990s versus the 1980s being more than offset with losses in the 2000s, which matches the model predictions. Even more striking, are the patterns for the college workers. For this group, both the 1980s and 1990s show a clear pattern of movement out of middle-and low-paying occupations and into high-paying occupations. After 2000, the pattern reverses, with declines in top-paying occupations and growth in the lowest-paying sectors. This fits with our model of a long term growth in employment in cognitive occupations offset by a reversal

---

more complex math skills are required for an occupation, and `DCP`, which measures adaptability or responsibility of decision, control and planning of an activity. Our cognitive task measure is $cog\_task = (\texttt{GED} + \texttt{DCP})/2$ at the disaggregated occupation level. When we compute this measure for more aggregated occupation categories, we take a weighted average of this measure where the weights are the proportion of hours worked in each occupation. Similarly, we calculate a measure of routine task intensity by using the codes `STS` and `finger`, which measure the extent that an occupation requires the use of "set limits, standards or tolerances" and "finger dexterity", and are argued to capture occupational activities which are routine and can be replaced by computing technology. Finally, we use the measure `EHF` which measures "Eye, hand, and foot coordination" as a measure for manual task intensity. Cognitive task occupations we use are based on Acemoglu and Autor (2011) and include Management, Professional and Technical occupations, and, as a group, have a *cog_task* measure is larger than the routine or manual occupations. More detail is given in the Data Appendx.

of this trend after 2000.

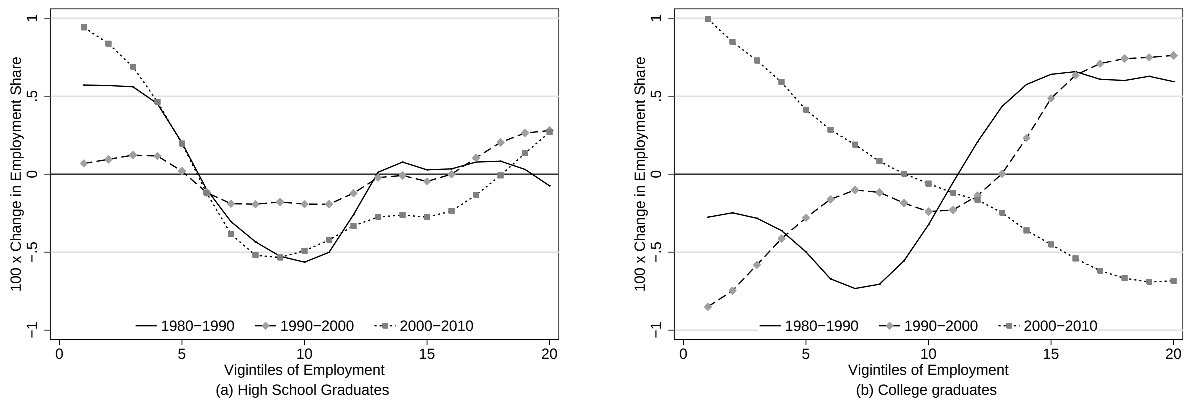

Fig. 11.— Smoothed Employment Changes by Skill Quartile: Young Workers

We provide further evidence on the occupation and wage experiences of cohorts of new college graduates in our paper, Beaudry et al. (2014). There, we show that the proportion of new graduates starting out in manual and professional jobs grew by over 10 percentage points between the early 1990s and 2000. But this was completely reversed by the time the 2010 graduates enter the labor market. For all cohorts, one can observe a pattern where at least a third of new graduates start out in manual or routine jobs and then a portion of them move into cognitive task occupations with time. The rate at which this latter movement occurred slowed down substantially in the 2000's relative to the 1990s. Thus, new college graduates are have become increasingly less likely to start out in a management or professional job and more likely to be trapped in service or routine jobs.

These patterns are reminiscent of the literature on the effects of graduating in a recession. Using Canadian data, Oreopoulos (2012) show that university educated workers who graduate in a recession have lower initial earnings and take approximately 10 years to catch back up to the earnings of luckier cohorts that do not graduate in a recession. Similarly, Kahn (2010) and Altonji et al. (2013) show significant effects of graduating in a recession in the US, with some of the initial difficulties being related to starting out in lower paying/lower prestige occupations. The patterns that we document can similarly be construed as reflecting reduced access to higher paying occupations during periods of reduced demand for labor. This brings us back to the earlier question of whether the reduction in labor demand after 2000 can be seen as essentially cyclical in nature versus reflecting a structural shift. We view results such as those in Figure 11 for college workers, in which the pattern of changes in the occupational density in the 1980s and 1990s is completely reversed after 2000, as fitting with the structural shift view, though we obviously cannot prove this. Only time will tell whether the reduced demand for labor continues or reverses once again.

Together, we view Figures 10 and 11 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. Wage Patterns

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 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 these patterns should properly be seen as referring to task prices for the three tasks emphasized in the model: cognitive, routine, and manual. 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 12 we plot two alternative measures of the average wage paid in cognitive (panel (a)), routine (panel (b)) and manual (panel (c)) 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.

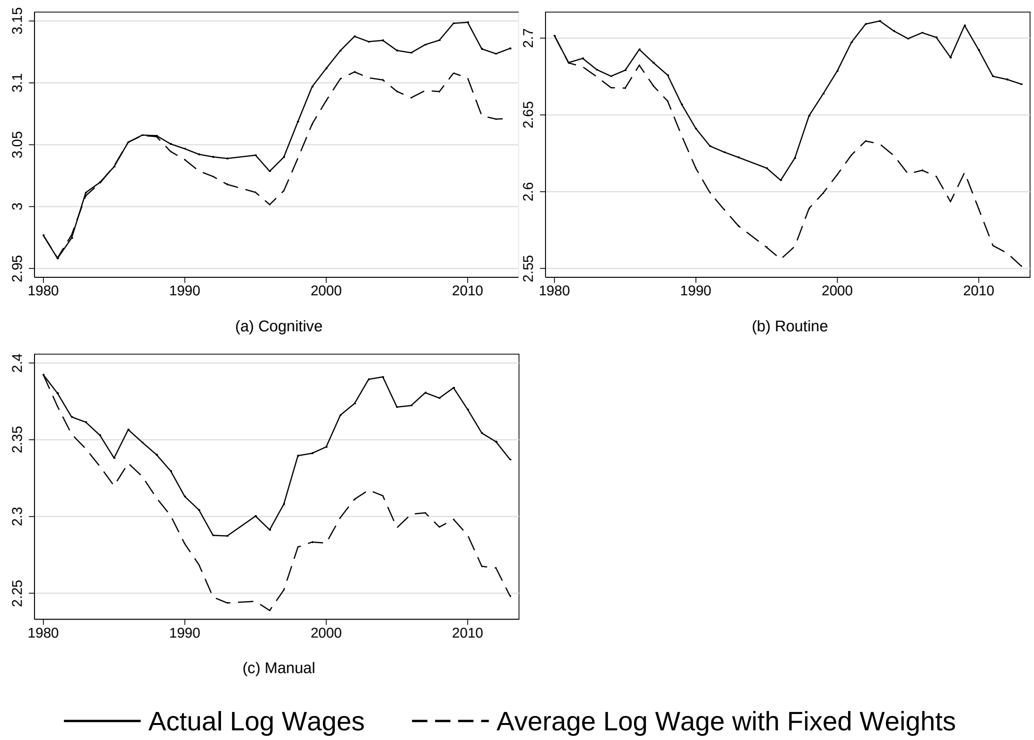

Fig. 12.— Average Wage by Occupation: Actual and Fixed Weight

We take two points from these figures. First, the cognitive wage follows a pattern much like cognitive employment, with strong growth in the 1990s followed by little or no growth after 2000. Similarly, routine and manual wages stop growing and start declining in the early 2000's. These wage patterns emerge well before the 2008 recession. Second, the fact that the difference between the observed wage lines and the composition constant wage lines increased after 2000 indicates that the skill composition in terms of observable measures (experience and education) increased within each occupation. This fits with the earlier observations that high school workers moved out of cognitive task jobs and college workers moved down the occupation ladder after 2000. If, as seems plausible, selection on unobservables moved in the same direction as the observable characteristics change then the implication is that the actual task prices declined after 2000. This would fit with the model implications that decreased demand for cognitive occupations drove the cognitive task price down after 2000 and increased supply drove down the manual task price.

In Figure z, we present differences in wages at each percentile of the wage distribution for both young (age 25 to 34) and older (age 35 and above) workers, broken down by education. If we are willing to assume that workers maintain their relative rank in the wage distribution within their age x education group over time then the shifts in percentiles in these figures are a step closer to controlling for both unobservable and observable skills. In that context, it is interesting to note that for young workers in both education groups, the 2000s correspond to a period of wage decline (with the declines being larger if we include the post-2008 recession years) while the 1990s wage changes were positive or zero across all percentiles. For older workers, the 2000s again corresponds to a shift down in the wage change line relative to the 1990s line at all percentiles, though for older college workers the changes are not negative across most percentiles. We view the plots for younger workers as more closely corresponding to the types of wage movements that allow us to understand the impacts of technology changes since wage patterns for older workers are more affected by implicit contracts, etc.. Given that, we view this as evidence that the key implications of the model in terms of declining real wages (once one controls for composition shifts) are present in the data. However, we recognize that a more complete investigation of selection effects would be useful and leave that for future work.

3.4. Corroborating Evidence and Competing Theories

To this point, we have presented evidence of a shift in the US labor market beginning around 2000 and set out a model for why that might have occurred that relies on the notion of a maturing General Purpose Technology. We see patterns in wages and employment within education groups as fitting with that theory but readers might reasonably demand more corroborating evidence. In this section, we present two types of corroborating evidence based on sector and on geography. We then turn to considering alternative theories for why the labor market changed in 2000.

3.4.1. Sectoral Patterns

One of the key pieces of evidence in favor of IT related technological change as a main driving force for US labor market patterns is the IT investment series presented earlier. Those series show trend breaks in 2000 in both equipment and software, with strong upward trends giving way to either declining or flat investment. In Figure 13, we show the ratio of investment in IT equipment and software to output for what we call IT intensive industries (industries where the ratio is over 5 at some point).[19] Apart from transportation, the IT intensive industries tend to be ones where we expect cognitive tasks to be important (Professional and Scientific, Education, etc.). For the information and cultural sector, the series rises sharply to its peak in 2000 then falls substantially. The Professional, Scientific, and Math sector also peaks in 2000. For other sectors, the peaks are more scattered, with Transportation peaking in 1997, Education peaking in 2002, and Finance having a local peak in 1999. But regardless of the exact timing of the peaks, these sectors all show patterns of flat or negative growth in the IT investment ratio in the 2000s. The only exception is the Management sector within which the ratio continues to grow throughout the period. Interestingly,

[19] In the figure, the ratio is multiplied by 100.

the pattern of relative stagnation of IT investment in the 2000s is present even in low IT intensive sectors such as Construction and Manufacturing. We plot the IT investment - output ratios for a set of these industries in Figure 14, where one can see a fall in the investment ratio in all sectors that has varying timing but in all cases occurs before the onset of the 2008 recession. We view this as reflecting IT being a General Purpose Technology (GPT) that transformed all sectors of the economy. All sectors, whether heavy industry or scientific services, invested in IT before 2000 and reduced their investment in it after the late 1990s.

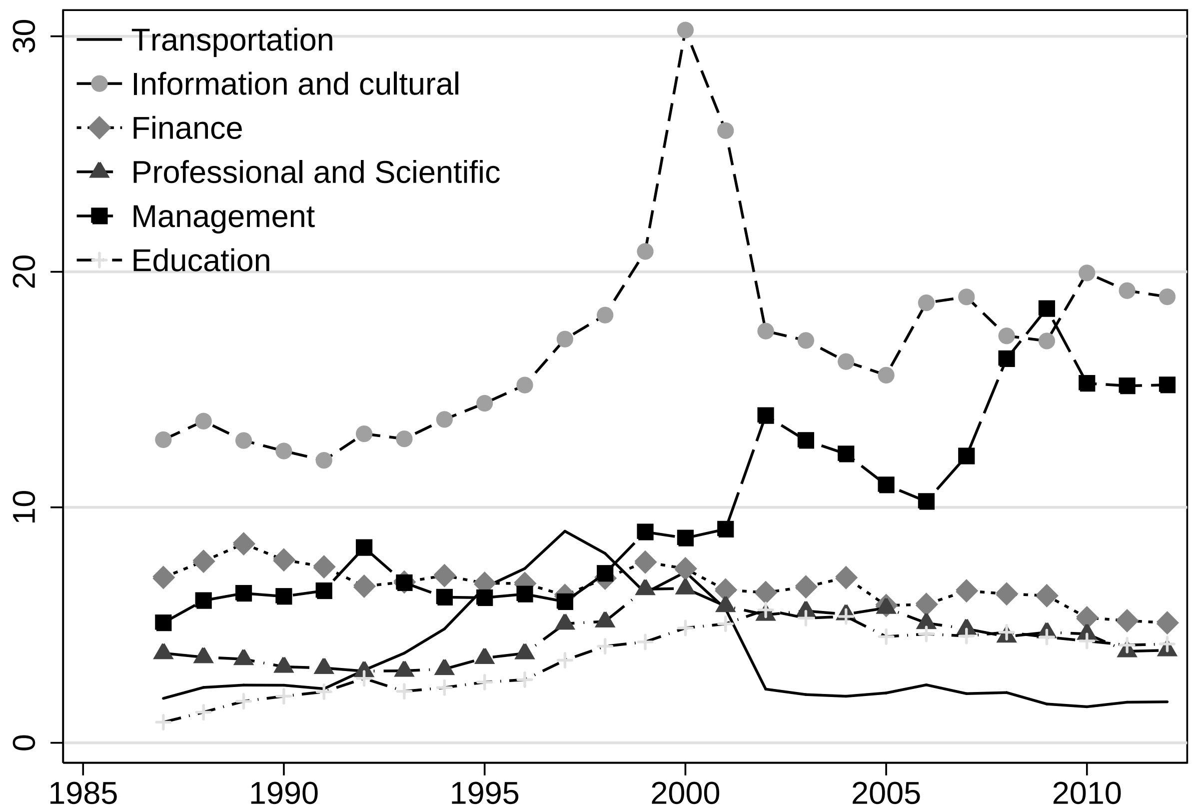

Fig. 13.— IT Investment to Output for IT Intensive Industries

IT being a GPT makes it more difficult to test the predictions of our model since impacts show up in all sectors, leaving us with variation in the form of a single time series. The same can be seen in employment data when we examine the rates of employment of cognitive task occupations by industry during the 1990s and 2000s. We perform a standard decomposition exercise in which we decompose changes in the share of employment in cognitive task occupations into between and

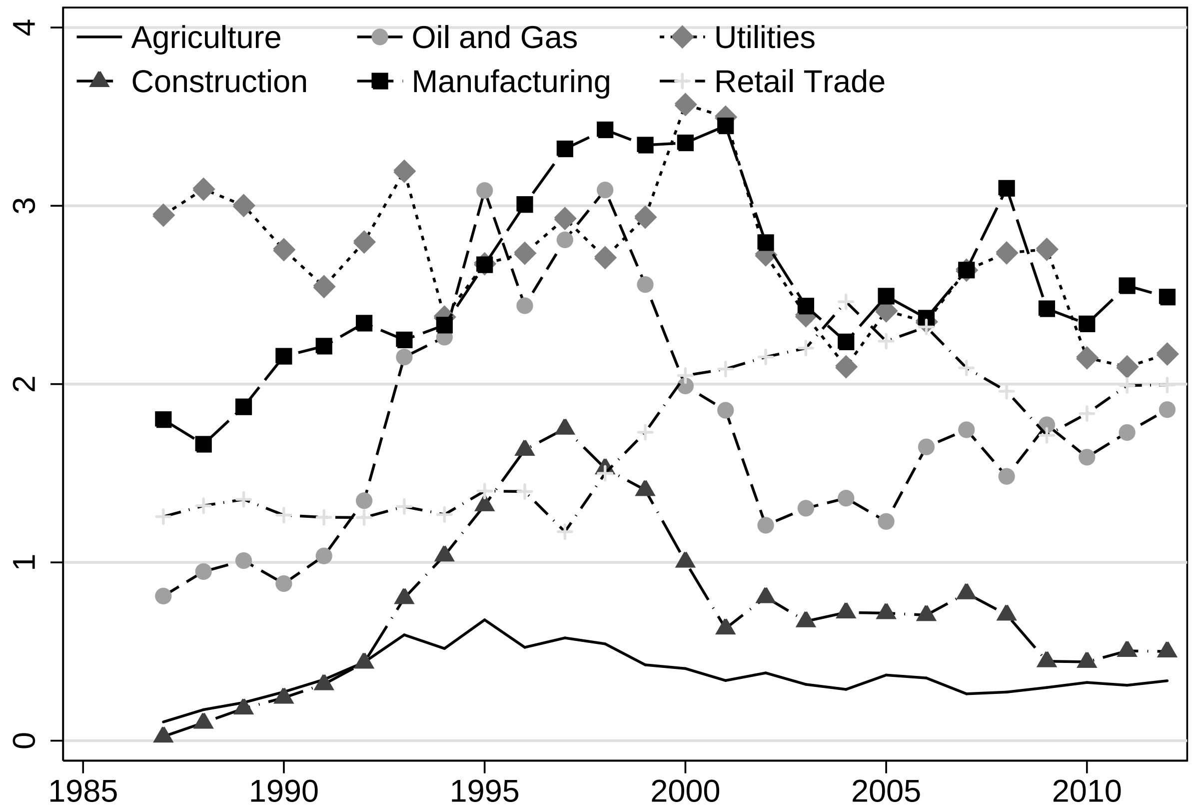

Fig. 14.— IT Investment to Output for non-IT Intensive Industries

within industry components.[20] For the 1990s, the overall share in cognitive task occupations for young (age 25 to 34) BA or more educated workers increased by 3.32, with 77.1% of the increase associated with the within component and the remainder accounted for by shifts in the industrial composition of employment for this group. For the 2000s, the overall share fell by 3.11, with 65% accounted for by the within component. Thus, the majority of the total change in cognitive task employment for young, well-educated workers happens within all industries, fitting with what we would predict for a GPT.[21]

[20]We use 15 two digit industries: agriculture; construction; manufacturing, durable goods; manufacturing, non-durable goods; transportation; communications; wholesale trade; retail trade; finance, insurance, and real estate; business and repair services; non-business services; hospitals; education services; social services; and public administration.

[21]For college workers of all age groups, the comparable numbers are: an increase of 1.34 in the 1990s, with the within component accounting for 140% of the change; and a fall of 3.0 in the 2000s, with 51% being due to the within component. Thus, the within component being the main determinant is also present for the overall college educated

There is one industry that holds specific interest even within the context of a GPT: business and repair services. This sector includes: research, development, and testing services; personnel supply services; management and public relations services; and computer and data processing services. If IT represents a GPT then one would expect firms in various industries to look to hire help in implementing the new technology and the workplace changes that go with it. The categories in this industry indicate that it contains the types of firms that others would turn to for this help. During the 1990s, the share of young college graduates in the business and repair services industry nearly doubled, from 17% in 1990 to approximately 30% in 2000, before declining to approximately 24% by the end of the 2000s. In Figure 15, we present a counterfactual exercise highlighting the importance of this sector for the overall decline in the proportion in cognitive task jobs among young college graduates. The solid line in the figure shows the actual cognitive proportion for this group.[22] The dotted line above it corresponds to a counterfactual exercise in which we held the proportion of workers within the management services industry at its 2000 level. Removing this industry from the overall within-industry component in this way reduces the 2000's decline in the cognitive share. For the next line up (the dashed line), we projected what the overall cognitive share would have been for young college educated workers if the proportion of them working in the management services industry had continued to grow as it did in the 1990s. The top line combines the latter between-industry counterfactual with holding the within-management-services cognitive proportion constant at its 2000 level. Under the latter counterfactual, the overall proportion in cognitive task occupations would not have declined in the 2000s. Thus, it appears that changes in the proportion working in the sector that supports the implementation of the IT GPT along with changes in the use of cognitive tasks within that sector played an important role in the turnaround in the overall demand for cognitive tasks after 2000. This fits with our notion that there was investment in putting in place systems that made use of the IT GPT before 2000 that slowed after 2000.

3.4.2. Cross-City Evidence

We next look for further 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

group, though it is more important in the 1990s but somewhat less important in the 2000s than for younger college educated workers. We view the young workers as the most useful for capturing newly emerging patterns because they are not yet subject to rigidities associated with longer term contracts and work arrangements.

[22] The original "business and repair services" industry also includes industries such as automotive repair. The counterfactual exercise in the figure is based on a constructed sector including only research, development, and testing services; personnel supply services; management and public relations services; and computer and data processing services. Dropping automotive repair and other similar sub-industries has little impact on our exercise because very few college graduates are employed in these categories.

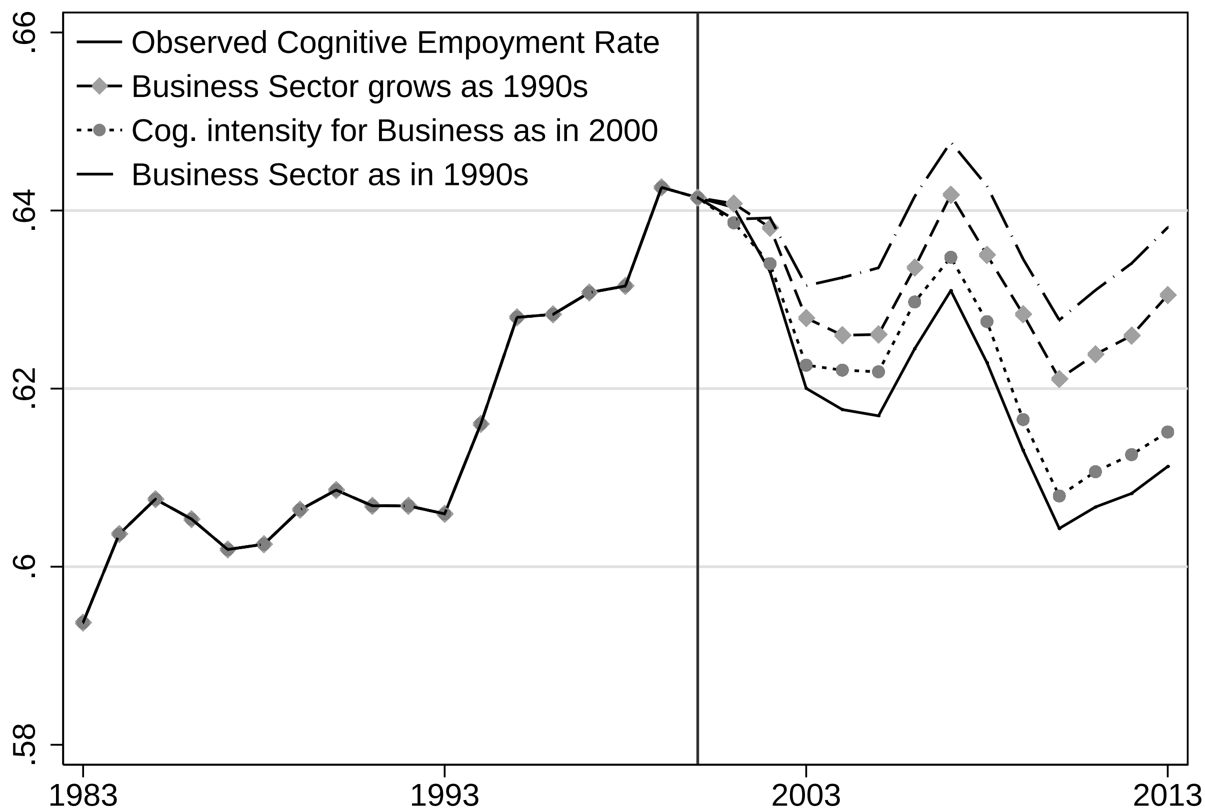

Fig. 15.— Employment Rate in Cognitive Occupations for Young College Graduates

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 θ (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 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:

$$\Delta ER_{c,2010-2000} = \alpha_0 + \alpha_1 X_{c,pre-2000} + \epsilon_{ct}, \tag{5}$$

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.

Table 1: Estimates of Equation (5)

	(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.

Table 1 presents the results of OLS estimates of equation (5). 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 1 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.

3.4.3. Competing Explanations

The turmoil in the US labor market in the 2000s has not, of course, gone unnoticed by other researchers, and other explanations have been advanced. Perhaps chief among these is the argument that the increased export capacity of China has had an important impact on the American economy. Autor et al. (2013) support this conclusion for employment and wages in US manufacturing using cross-regional variation in the predicted impact of Chinese good import penetration. At the aggregate level, they show that the long term decline in employment in manufacturing accelerated after China joined the WTO in 2001, at which point the growth in the import penetration ratio of goods from China accelerated upward. The timing of this shift is quite similar to the change we have documented throughout the US labor market in 2000. We see the relationship between these changes as nuanced. On one side, we have just seen that the changes in the proportion working in cognitive task jobs after 2000 is spread across all sectors, not just manufacturing. Moreover, it seems unlikely to us that increased imports of goods produced using low skilled labor in China would be a direct determinant of reduced demand for cognitive task occupations. On the other side, it does seem plausible that at least part of what we see in terms of reductions in routine task employment - and, particularly, its acceleration downward after 2000 – could be stemming from trade with China as well as from technological change. In that case, the apparent increase in supply of workers to the manual/service sector after 2000 could also be partly due to trade effects. At the same time, the increased outsourcing of jobs typically held by those with a BA, particularly to India, could imply that some of what we observe in terms of cognitive task trends is related to trade (though, possibly, not to China). Because this type of outsourcing is only possible because of the IT revolution, we would view such effects as potentially part of what we are describing.

A second possibility is that the post-2000 period represents a new technological era: one where technologies have turned what people who have exactly a BA do into more routine tasks that can be replaced with technology or outsourced. At the same time, demand for those with more advanced degrees could have increased. In other words, the same notion of technological change as a polarizing force could have continued but what occupations are in the middle verus upper part of the polarization has shifted upward. Lindley and Machin (2013) document the increased share of employment accounted for by advanced degrees. While this may be part of what has been happening, we don't believe it is the main story for two reasons. First, all the patterns we have presented work with a combination of BA's and those with advanced degrees. When we work only with those with exactly a BA, we observe very similar patterns. Similarly, in Beaudry et al(2014), we find similar growth and decline in the proportion in cognitive task employment among the exact BA's and those with advanced degrees. Second, to the extent the advanced degree holders are following a different trend, we would expect to see a different employment pattern in the very high wage occupations in Figure 9. Instead, we see a decline in the growth rate in employment in all top paying occupations.

A third possibility is that there was a general, protracted downturn in the US economy beginning around 2000, and that showed up in college educated workers moving down through the occupational structure. It is not clear that we can identify against a broad story of downturn in the labor market driven by unspecified forces, but several pieces of evidence incline us toward our story of stagnation associated with a maturing GPT. First, while it is true that employment in cognitive task jobs stalls or declines slightly for college educated workers in the 1980s (as fits with work on effects of graduating in a recession by Kahn (2010) and Oreopoulos (2012)), these appear as relatively small interruptions in the longer term pattern of growth. In contrast, the break downward in the proportion in these jobs among college graduates is stark and large. Second, investment in IT related capital (both equipment and software) also shows a significant break in 2000. Third, the management services industry shows a particularly strong pattern of increased employment in the 1990s followed by decline after 2000 and played a substantial role in the overall cognitive employment pattern.

3.5. 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 6 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. 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 crisis likely play a role, and the structural decline in employment in routine occupations and manufacturing jobs are certainly contributing factors (Charles et al. 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. While this demand for cognitive tasks directly affects 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 task occupations, 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 for the changes pre- and post-2000.

Finally, it is interesting to consider what this hypothesis might imply for the recovery of the US labor market in the aftermath of the Great Recession. For those who see the recession as a cyclical demand deficiency problem, the labor market can make a full recovery once sufficient stimulus is applied to the economy. For those who blame persistent employment problems on structural changes on the supply side of the labor market, stimulus is not the answer. Instead, they advocate making the labor market more flexible by, for example, reducing unemployment insurance generosity. Our argument is that the problems in the labor market pre-date the recession and that while they are structural in nature, the relevant structural shift is on the demand side rather than the supply side of the market. In this sense, our results fit with the work of (Charles et al. 2012) and (Siu and Jaimovich 2012) who argue that long term problems in the manufacturing sector were temporarily masked by the early 2000s housing boom. Like them, we would argue against expecting a return to persistently strong demand for labor in the near future even if governments either increased stimulus or cut benefits. In contrast to them, we look beyond the manufacturing sector to the whole labor market and lay the blame on patterns of technological change that have implied a reversal in the demand for the more educated. In that sense, the direct implications of our argument are more pessimistic: the US cannot expect good demand for labor even at the top end of the labor market in the near term. More optimism awaits a new technological revolution.

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 4*, 1043–1171.

Altonji, J. G., L. B. Kahn, and J. D. Speer (2013). Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success. Technical report, Yale University.

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, August). 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., D. Dorn, and G. H. Hanson (2013, October). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review 103*(6), 2121–2168.

Autor, D. H., L. F. Katz, and M. S. Kearney (2006). The Polarization of the U.S. Labor Market. *The American Economic Review, Papers and Proceedings 96*(2), 189–194.

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

Bartel, A. P. and F. R. Lichtenberg (1987). The Comparative Advantage of Educated Workers in Implementing New Technology. *The Review of Economics and Statistics 69*(1), 1–11.

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

Beaudry, P., D. A. Green, and B. M. Sand (2014, May). The Declining Fortunes of the Young Since 2000. *American Economic Review: Papers & Proceedings 104*(5), 381–386.

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. Technical report, University of Chicago.

Dustmann, C., J. Ludsteck, and U. Schönberg (2009, May). 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. Technical Report 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.

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

Kahn, L. B. (2010, April). The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics 17*(2), 303–316.

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.

Lindley, J. and S. Machin (2013). The Rising Postgraduate Wage Premium. Technical Report October.

Oreopoulos, P. (2012). The short-and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics 4*(1), 1–29.

Schultz, T. (1975). The Value of the Ability to Deal with Disequilibria. *Journal of economic literature 13*(3), 827–846.

Siu, H. E. and N. Jaimovich (2012). The Trend is the Cycle: Job Polarization and Jobless Recoveries. Technical report, NBER Working Paper No. 18334.

Welch, F. (1970). Education in Production. *The Journal of Political Economy 78*(1), 35–59.
