

Labor Market Polarization, the Decline of Routine Work, and Technological Change: A Quantitative Analysis^{*}

Christian vom Lehn[†]

This Draft: December 22, 2014

Abstract

Technological change is a prominent hypothesis for the recent polarization of the labor market and the related decline for occupations specializing in performing routine tasks. In this paper, I provide a quantitative evaluation of this hypothesis. To do so, I build an extension of the standard growth model which allows for endogenous determination of the demand and supply for occupational labor in response to investment specific technological change. I further evaluate the extent to which this channel of technological change can account for recent declines in aggregate employment and the labor share of income. My analysis finds that technological change is able to account for a large fraction of changes in occupational employment and earnings, as well as the decline in the labor share, through the year 2000, but is unable to reconcile many of these patterns in the subsequent decade. In particular, after 2000, the model significantly overpredicts wages and hours for higher skilled occupations. This is at odds with both the recent measured slowdown in demand for these occupations as well as the hypothesis that slowing technological change can account for this phenomenon.

JEL codes: E24, J24, O33

^{*}I am deeply indebted to Richard Rogerson, Esteban Rossi-Hansberg, Mark Watson, and Tom Winberry for their feedback on this project. This paper is a substantially revised version of chapter 1 of my PhD dissertation completed at Princeton University. I am also very grateful for helpful comments and suggestions from Mark Aguiar, David Autor, Jesus Fernandez-Villaverde, Mike Golosov, Oleg Itskhoki, Greg Kaplan, Nobu Kiyotaki, Val Lambson, Alex Mas, Ezra Oberfield, Pedro Olea, Ricardo Reyes-Heroles, Chris Sims, Gabriel Tenorio, Sharon Traiberman, Jon Vogel, and participants in seminars at the Princeton Macro Workshop, Brigham Young University, Utah State University and the TASKS III conference. I am also grateful to Riccardo DiCecio for sharing his data for the quality-adjusted price of investment goods.

[†]Brigham Young University. E-Mail: cvomlehn@byu.edu.

1 Introduction

The work of Autor et al. (2006, 2008), has focused attention on the polarization of the labor market - the comparatively slow growth in employment and wages for middle-skilled occupations relative to lower and higher skilled occupations since the 1980s.¹ Because these middle-skilled occupations specialize in performing routine tasks, a prominent hypothesis for this polarization, originally set forth in Autor et al. (2003), is the rise of routine-biased technical change. Persuasive suggestive evidence in favor of this hypothesis has been presented recently in Michaels et al. (2013), Autor and Dorn (2013) and Goos et al. (2014). However, the aggregate effects of routine-biased technical change have yet to be evaluated quantitatively in general equilibrium. In this paper, I conduct a formal quantitative evaluation of this hypothesis by developing a dynamic general equilibrium model of occupational labor supply and demand.

My model augments the standard growth model by adding three key features: (1) the occupational tasks framework of Autor and Dorn (2013), (2) a model of occupational choice across heterogeneous worker skill levels, patterned after Eaton and Kortum (2002) and Cortes (2014), and (3) equipment investment specific technical change, as in Greenwood et al. (1997). Because I distinguish between the skills of individual workers and the task content of occupations, my framework is consistent with new evidence I present that middle-skilled *workers* have not experienced negative labor market outcomes comparable to those of middle-skilled *occupations*. The quantitative evaluation I conduct is accomplished by feeding in a measured series of equipment investment technological change and observing the degree to which the model can account for the time paths of shares of total hours worked and average hourly earnings across occupations.

My analysis generates two key findings. First, I find that technological change accounts for much of the observed labor market polarization through the year 2000. In particular, the model accounts for virtually all of the changes in occupational hours shares and a significant fraction of changes in relative hourly earnings across occupations. Second, I find that technological change is much less able to adequately account for polarization since 2000. In particular, the model sub-

¹See Acemoglu and Autor (2011) for a recent survey of the literature that has followed.

stantially overpredicts the growth of hourly earnings and the share of hours worked in the highest skilled occupations.

My model also allows me to evaluate the extent to which this form of technological change can account for other important changes in the aggregate economy, such as recent declines in the labor share of income and the aggregate employment to population ratio.² The ability of technological change to account for these patterns closely mirrors the above results regarding polarization. Specifically, the model is consistent with the decline in the labor share and the level of the employment rate through 2000, but the model is unable to account for subsequent declines in these series. A decomposition exercise highlights that a significant reason the model is unable to account for recent declines in the employment to population rate and the labor share is because it overstates of the share of income paid to higher skilled occupations and the employment rate of workers in these occupations.

This last finding is related to work by Beaudry et al. (2013), which documents a decline in demand for higher skilled labor since the year 2000. They suggest that this decline in demand may be due to a slowdown in technological change, but my results show that the slowdown in technological change is insufficient to account for this pattern. In particular, the model accounts for the recent slowing in technological growth, but is unable to reconcile the slowing growth in higher skilled wages and hours, suggesting the need for an alternate explanation for these recent trends.

This paper relates to a large literature that quantitatively evaluates the impact of technological change on shaping labor market outcomes, including Krusell et al. (2000), He and Liu (2008), Burstein et al. (2014), and Morin (2014). While Krusell et al. (2000) and He and Liu (2008) focus on the skill premium, Burstein et al. (2014) and Morin (2014) focus more directly on the interaction of tasks and technological change. Burstein et al. (2014) uses a task structure to understand changes in wages across worker skill groups, and find that technological change plays a substantial role in shaping these changes. Relative to Burstein et al. (2014), I focus primarily on occupational,

²Recent studies by Karabarbounis and Neiman (2014) on the labor share and Jaimovich and Siu (2012) on the aggregate employment rate have similarly considered such connections.

or task-related outcomes, and evaluate the role of technology in general equilibrium, focusing on both wages and hours, and not just the former. Further, their exercises only go through the year 2003, whereas my quantitative exercise goes through the year 2012 and highlights important differences in the decade of the 2000s with regard to technological change and labor market outcomes which are not evident in their work.

Morin (2014) considers the extent to which routine-biased technical change can account for recent sluggish cyclical employment recoveries as well as the trend decline in labor share. In contrast to my results, Morin (2014) suggests that routine biased technical change can potentially account for the entire labor share decline. My model framework similarly admits the same theoretical possibility, but in calibrating directly to micro data and using an actual measured series for technological change, I obtain the different result that technological change is unable account for this decline after 2000.

The remainder of the paper is structured as follows. Section 2 gives background on labor market polarization and its relationship with worker skills and occupational tasks, then documents the data facts to be considered in the quantitative evaluation. Section 3 lays out the model framework. Section 4 describes the measurement of technological change in the data and the solution procedure for the model, and Section 5 presents the calibration strategy. Section 6 presents the aggregate results from feeding in the data series on technological change, and Section 7 provides some detailed individual level analysis of those results. Section 8 concludes.

2 Polarization and Aggregate Labor Market Changes

2.1 Polarization, Skills, and Tasks: Background

2.1.1 Polarization and Skills

Following Autor and Dorn (2013), Figure 1 plots the log changes in real wages employment shares across the occupational skill distribution between 1982 and 2012.³ Employment and wages have grown most rapidly for occupations in the lowest and highest skill percentiles, and have grown much slower for occupations in roughly the middle 60% of the skill distribution. This is what has been termed the polarization of the labor market.

An important observation regarding labor market polarization is that it focuses on labor market outcomes for given *occupations*, not workers. Thus, it is unclear whether the same negative labor market outcomes have been observed for middle-skilled workers, or rather, the workers typically employed in middle-skilled occupations, as these workers may instead choose to work in higher or lower skilled occupations.

I present evidence of differences in the labor market outcomes for middle-skilled workers and middle-skilled occupations by the employment outcome of workers traditionally employed in middle-skilled occupations. To do this, I define the employment rate for an occupation j as:

$$\left(\frac{E}{P}\right)_{jt} = \sum_s \frac{E_{jst}}{E_{jt}} \frac{E_{st}}{P_{st}} \quad (1)$$

where E_{jt} represents the amount of employment in occupation j at time t , E_{jst} represents the amount of employment in occupation j by workers of skill s , E_{st} represents total employment of workers of skill s , and P_{st} represents the population of individuals with skill s . In short, the occupational employment rate is the weighted average of the employment rate for each of the

³Occupations are ranked by their mean log wage in 1982 and divided into 100 percentiles based on this ranking, forming the mentioned “occupational skill distribution.” This procedure follows the same general construction methods for these figures as in the literature at large, particularly Acemoglu and Autor (2011) and Autor and Dorn (2013). Results are robust to excluding the recession of 2007-2009 and considering the growth only through the year 2007. Additional details regarding data sources and construction can be found in Appendix A.

worker skill levels employed in that occupation.

Given this definition, I ask the question - what has happened to the employment rates for the types of workers who traditionally chose employment in each occupation? That is, I consider the growth in occupational employment rates, with the weights in (1) , $\frac{E_{jst}}{E_{jt}}$, fixed at their level in the year 1980.⁴ This allows me to identify any differences between worker and occupational outcomes. If middle-skilled workers are seeing similar labor market outcomes as middle-skilled occupations, one might expect growth in these employment rates to also exhibit a U-shape pattern, as in Figure 1.⁵

Figure 2 shows the results of this exercise. Although employment growth has been U-shaped across the occupational skill distribution from 1982-2012 , employment rate growth over the same span has been monotonically increasing across the same distribution.⁶ While this result does not have clear implications for the exact changes in occupational choice patterns by middle-skilled workers, it is consistent with recent studies on occupational switching, which have suggested that workers in middle-skill occupations are switching to both higher or lower skilled occupations (and potentially displacing lower skilled workers from their occupations).⁷ In particular, this result suggests that a complete understanding of labor market polarization not only encompasses understanding why there has been decreased demand for certain types of occupations, but also how workers have changed their occupational choice patterns in response.

2.1.2 Polarization and Tasks

A common framework that has emerged from the literature on polarization identifies two broad groups of tasks performed by workers - routine tasks and non-routine tasks. A routine task is

⁴To do this, I group the population into 45 skill categories - 5 educational categories and 9 potential experience categories - and compute the log difference in occupational employment rates between 1982 and 2012, with the weights fixed at the year 1980. Further details are available in the notes to Figure 2.

⁵Technically, if there were also substantial changes in skill populations, this might not be the case. However, accounting for such changes does not overturn the results of the exercise.

⁶This result is robust to a variety of measurement approaches, including measuring employment rates as total hours worked in the prior year per population, the employment rate at a given point in time, or even as measuring total employment (which accounts for population growth).

⁷Examples of such studies include Cortes (2014) and Cortes et al. (2014).

highly repetitive, follows a strict set of rules or procedures, and requires precision and accuracy. On the other hand, non-routine tasks require either creativity, abstract reasoning, persuasion or some degree of physical exertion and dexterity or direct interpersonal interaction. While different studies have further decomposed these two broad groups of tasks in different ways, I follow Autor and Dorn (2013) and consider three types of tasks - abstract, routine and manual tasks.⁸ Abstract tasks are those non-routine tasks which require intuition, creativity, persuasion and/or abstract reasoning. Manual tasks are those non-routine tasks which require physical exertion, strength and/or direct interpersonal interaction.

Although most jobs perform each of these three types of tasks to some extent, occupations can be sorted by the tasks which are most prominent in the work they do. Autor and Dorn (2013) use data from the Dictionary of Occupational Titles to identify different task intensities across occupations and classify occupations according to these three task groups; I use these same classifications.⁹ Examples of occupations that specialize in each of these three types of tasks are:

Abstract: teachers, doctors, lawyers, engineers, economists

Routine: secretaries, bookkeepers, retail salespersons, machine operators, assembly line workers

Manual: hairdressers, auto mechanics, bus drivers, housekeepers, construction laborers

Both Acemoglu and Autor (2011) and Autor and Dorn (2013) highlight how occupations specializing in routine tasks tend to be more concentrated in the middle of the skill distribution, and thus are potentially closely related to the observed employment and wage polarization. Further, manual tasks are more concentrated in lower skilled occupations and abstract tasks are more concentrated in higher skill occupations. For the remainder of the paper, I focus my attention solely on labor market outcomes for these three types of occupations.

⁸The seminal work on tasks and labor market outcomes by Autor et al. (2003) originally introduced five categories of occupations; more recent work by the authors has focused on three or four categories, as in Autor and Dorn (2013) and Acemoglu and Autor (2011).

⁹Using alternative classification systems, however, such as those of Jaimovich and Siu (2012) lead to similar conclusions in the empirical analysis.

2.2 Labor Market Trend Dynamics

I present the sets of facts I will focus on in my quantitative analysis: the shares of hours worked and relative average hourly earnings across occupations, as well as the aggregate employment rate and labor share of income. For all these trends, I focus particularly on the time paths of these changes; this contrasts with much of the early literature on polarization, which has only focused on changes between two points in time.¹⁰ Data comes from the CPS March Supplement unless otherwise specified. More details regarding data construction are presented in Appendix A.

2.2.1 Polarization: Hours Shares and Relative Hourly Earnings Across Occupations

Figures 3 and 4 plot the shares of total hours across occupations and the average hourly earnings in routine occupations divided by those in abstract and manual occupations.¹¹ Since the early 1980s, the share of total hours worked in routine occupations has been steadily declining, and the real wage for these occupations has grown slower than (or at least as slow as) other occupations.¹² The majority of hours and wage growth over this time period has come in abstract occupations, though there has been a slight surge in hours worked in manual occupations observed in the decade of the 2000s.¹³ These trends are consistent with the evidence already seen on polarization across the broader spectrum of occupations.

2.2.2 Aggregate Patterns: Employment to Population Ratio and the Labor Share of Income

Figures 5 and 6 plot the employment to population ratio and the labor share of income for the total economy. Through the year 2000, there was an upward trend in the percentage of the population

¹⁰More recently, Foote and Ryan (2012), Smith (2013), and Cortes et al. (2014) have also given attention to these time paths.

¹¹I focus on the shares of hours worked as opposed to employment shares, as there can be substantial variation in the average number of hours worked across occupations.

¹²One common concern is that this decline, particularly in hours, is simply capturing changes in industry demand occurring with the decline of manufacturing. Acemoglu and Autor (2011) and Jaimovich and Siu (2012) show that such a between industry shift can account for no more than a third of the decline. Further, Goos et al. (2014) suggest that a theory of routine-biased technological can also account for a significant share of between industry shifts.

¹³This surge is unlikely to be completely explained by a boom in construction preceding the Great Recession, as some of the growth in the manual hours share has remained since the recession. Autor and Dorn (2013) argue that a significant component of this is the rise of low skill service occupations.

that was employed, and a slight, steady decline in the labor share.¹⁴ However, since the 2000s, there have been substantial declines in both these series.¹⁵ Ongoing discussions have considered whether these declines are due to purely cyclical factors, or if there are deeper structural forces at work as well.¹⁶

Novel perspectives on both of these aggregate series are obtained by examining income shares and employment rates at an occupational level. As in Section 2.1.1, it is possible to compute the employment rate by occupation for abstract, routine and manual occupations, holding the distribution of worker skill in a given occupation fixed at 1980. Figure 7 plots changes in these occupational employment rates from 1975-2012. And Figure 8 plots the occupational labor shares - total income paid to an occupation divided by nominal GDP - for abstract, routine and manual occupations.

For these occupational employment rates, the same pattern observed in Figure 2 also holds true for abstract, routine, and manual occupations - employment rate growth increases as the initial skill content of the occupation increases. And while the aggregate labor share has been declining, particularly in the 2000s, the shares of income paid to each occupations in Figure 8 have been changing quite a bit over time. Notably, the most stark change among labor shares in the 2000s, when the aggregate labor share appears to show an accelerated decline, was the slowdown of growth in the labor share for abstract occupations.

3 Model

3.1 Environment

Time in the model is discrete, and one period in the model corresponds to one year.

¹⁴Much of the early growth in the employment to population ratio can be attributed to a rise in female labor force participation. See Juhn and Potter (2006) for a discussion.

¹⁵Karabarbounis and Neiman (2014) have documented this ongoing decline for the labor share worldwide, though Elsby et al. (2013) argue that a modest fraction of the decline in the US is due to accounting procedures used by the BLS.

¹⁶See Jaimovich and Siu (2012) for a recent argument about how trend movements in job polarization could be embodied in these cyclical fluctuations in aggregate employment.

3.1.1 Technology and Firms

An aggregate production function produces output using five inputs - structures (K_{st}), equipment (K_{et}), and efficiency units of manual (N_{mt}), routine (N_{rt}), and abstract (N_{at}) labor:

$$Y_t = K_{st}^\alpha \left[\mu_m N_{mt}^\sigma + (1 - \mu_m) \left[(1 - \mu_a) \left[(1 - \mu_r) K_{et}^\gamma + \mu_r N_{rt}^\gamma \right]^{\frac{\rho}{\gamma}} + \mu_a N_{at}^\rho \right]^{\frac{\sigma}{\rho}} \right]^{\frac{(1-\alpha)}{\sigma}} \quad (2)$$

The terms in square brackets represent three levels of aggregation. The innermost square brackets combine routine labor and equipment, to produce routine tasks, with the elasticity of substitution between routine labor and equipment given by $\frac{1}{1-\gamma}$. Routine tasks are then combined with abstract labor, with an elasticity of substitution given by $\frac{1}{1-\rho}$. Finally, the combination of routine and abstract tasks is combined with manual labor, with an elasticity of substitution given by $\frac{1}{1-\sigma}$. This production technology is a natural generalization of the tasks framework put forth in Autor and Dorn (2013).¹⁷

Output can be used for consumption or investment in structures or equipment according to the aggregate resource constraint:

$$Y_t = C_t + I_{st} + I_{et}$$

Following Greenwood et al. (1997), I assume capital accumulates as follows:

$$K_{st+1} = (1 - \delta_s) K_{st} + I_{st}$$

$$K_{et+1} = (1 - \delta_e) K_{et} + q_t I_{et}$$

¹⁷Relative to Autor and Dorn, I add structures as an input to production and do not impose $\rho = 1$. Further, Autor and Dorn (2013) consider two output sectors, goods and services, where services is only produced by manual labor and goods are produced by a combination of routine and abstract task inputs. Because household utility is a CES aggregate of both outputs, the final consumption good in their model is the same as mine.

where q_t represents equipment investment specific technological change. Movement in q_t is exogenous and deterministic.

With these laws of motion for capital accumulation, the aggregate resource constraint can be rewritten as:

$$Y_t = C_t + K_{st+1} - (1 - \delta_s)K_{st} + \frac{1}{q_t}(K_{et+1} - (1 - \delta_e)K_{et}) \quad (3)$$

3.1.2 Household

There is a household comprised of a continuum of individuals. Individuals are identified by a vector $\mathbf{s} = (z, z^a, z^r, z^m, \chi)$, where z is an individual's general skill, z^j is an individual's occupation-specific skill for occupation j , and χ represents the disutility an individual incurs from working. General skill, z , is distributed continuously across the real line, with $z \sim F(z)$, each occupation-specific skill, z^j , is distributed Fréchet with shape parameter θ , and χ follows a uniform distribution, $\chi \sim UNIF(0, \bar{\chi})$.¹⁸ The distributions over general and occupation-specific skill as well as disutility of work are independent of each other, with joint distribution $s \sim \Gamma(s)$, which is constant throughout time.¹⁹

Each individual in the household has preferences over consumption and labor supply, with period flow utility for an agent i given by:

$$U(C(s_i), E(s_i)) = \log(C(s_i)) - E(s_i)\chi_i$$

where $C(s_i)$ represents consumption by agent i and $E(s_i)$ is a 0/1 choice regarding employment.²⁰

Workers can be employed in one of three types of occupations - abstract, routine or manual.

¹⁸The choice of a Fréchet distribution is common in discrete choice settings; the choice of uniform distribution is for transparency and tractability. However, the results of the quantitative exercises are not meaningfully changed by using an alternate labor disutility distribution such as the Exponential distribution.

¹⁹Some discussion of whether allowing for a changing skill distribution over time would change the results of the paper is taken up at the end of Section (7.2).

²⁰Since the only variation in labor supply is on the extensive margin, this means that employment allocations and hours worked allocations will be identical in the model.

Conditional on employment in occupation j , $j \in \{a, r, m\}$, a worker with general skill level z and occupational skill z^j provides $z^j \phi_j(z)$ efficiency units in production, where $\phi_j(z)$ is an occupation specific productivity function, with $\phi_j(z) \geq 0$. As in Jung and Mercenier (2010) and Cortes (2014), I assume that workers with higher skill have a comparative advantage in abstract jobs relative to routine jobs and in routine jobs relative to manual jobs. Specifically, this is imposed by the following condition, akin to log supermodularity:

$$0 \leq \frac{d \ln(\phi_m(z))}{dz} < \frac{d \ln(\phi_r(z))}{dz} < \frac{d \ln(\phi_a(z))}{dz} \quad (4)$$

Earnings of a worker employed in occupation j are given by the product of the efficiency units they provide and the wage per efficiency unit in that occupation: $w_{jt} z^j \phi_j(z)$. Workers choose to work in the occupation that maximizes their income, given by $W_t(s) = \max_j w_{jt} z^j \phi_j(z)$.

The household takes prices as given and chooses consumption and employment for each individual, as well as next period's capital stocks to maximize an equal weighted integral of individual utilities, subject to their budget equation for each period:

$$\max_{C_t(s), E_t(s), K_{st+1}, K_{et+1}} \int_s U(C_t(s), E_t(s)) d\Gamma(s)$$

$$s.t. \quad \int_s C_t(s) d\Gamma(s) + K_{st+1} + (1 - \delta_s) K_{st} + \frac{1}{q_t} (K_{et+1} - (1 - \delta_e) K_{et}) = r_{et} K_{et} + r_{st} K_{st} + \int_s W_t(s) E_t(s) d\Gamma(s)$$

as well as subject to non-negativity constraints on capital and consumption and given initial conditions on capital stocks.

3.1.3 Firm

A representative firm produces aggregate output using the production technology in (2) and behaves competitively, renting equipment and structures from the household at rental rates, r_{et} and r_{st} , and paying wages w_{jt} per efficiency units of labor provided in occupation j , $j \in \{a, r, m\}$. The

firm's problem is standard and is given by:

$$\max_{K_{st}^D, K_{et}^D, N_{mt}, N_{rt}, N_{at}} \pi_t = Y_t - r_{st}K_{st}^D - r_{et}K_{et}^D - w_{mt}N_{mt} - w_{rt}N_{rt} - w_{at}N_{at}$$

3.2 Equilibrium

A competitive equilibrium in this model is an set of allocations $(C_t(s), E_t(s), K_{et}, K_{st}, K_{et}^D, K_{st}^D, N_{at}, N_{rt}, N_{mt})$, such that given prices $(r_{et}, r_{st}, w_{at}, w_{rt}, w_{mt})$:

- $(C_t(s), E_t(s), K_{et}, K_{st})$ solves the household problem, taking $(w_{mt}, w_{rt}, w_{at}, r_{st}, r_{et})$ and the initial conditions as given
- $(K_{et}^D, K_{st}^D, N_{mt}, N_{rt}, N_{at})$ solve the firm problem, taking $(w_{mt}, w_{rt}, w_{at}, r_{st}, r_{et})$ as given
- Labor, Capital and Output Markets clear

Appendix B reports the full set of equilibrium conditions for the model.

4 Measuring Technological Change and Solution Method

4.1 Measuring q_t

The model's sole exogenous driving force is equipment investment specific technological change, q_t . The equilibrium resource constraint in (3) implies that q_t is equivalently interpreted as the amount of investment that can be purchased with one unit of consumption - the relative price of equipment investment. Thus, as in Greenwood et al. (1997) and Krusell et al. (2000), I measure q_t as the price of consumption divided by the price of equipment investment. To measure the price of equipment investment, I utilize the equipment price series of DiCecio (2009), who extends through the year 2011 the quality bias price corrections to the BEA's investment price index performed originally by Cummins and Violante (2002) and Gordon (1990). To measure the price of consumption, I use BEA data on the price indices for non-durable consumption and non-housing

and non-energy services, and then apply quality bias adjustments to these series as suggested in Boskin et al. (1996) and Gordon (2006). Details of these bias adjustments are provided in Appendix A. The ratio of the consumption price divided by the equipment price yields a series for q_t through the year 2011, which I then smooth using an HP filter with parameter 6.25. Figure 9 shows q_t from 1980 to 2011.

A relevant question here is whether using the price of total equipment is the most appropriate measure in computing technological change. Remaining agnostic on which specific types of equipment are more or less substitutable for routine labor, I use total equipment as a benchmark for this analysis. As mentioned in Cummins and Violante (2002), growth in the technological change measure, q_t , from the early 1980s onward is substantially driven by both the increased share of investment in and the price decline of information and communications technology (ICT) equipment, suggesting that this benchmark may be appropriate relative to a narrower definition of equipment.²¹

4.2 Solution Method

Although the data series for q_t is only available through the year 2011, the model's results for the year 2011 and earlier depend on agents' beliefs regarding the time path for q_t beyond 2011. One possible approach is to allow q_t to continue to grow at some constant rate along a balanced growth path. However, as discussed in Greenwood et al. (1997) and He and Liu (2008), a balanced growth path will require production to be Cobb-Douglas so that q_t can be expressed as labor-augmenting technical change. As there are substantial secular trends in factor income shares as shown in Section 2.2.2, requiring production to be Cobb-Douglas is unappealing in this context.

Instead, I assume that equilibrium objects asymptotically converge to a final steady state where all variables are constant. Namely, I view the changes in q_t as inducing a transition of the economy from one steady state to another. I assume that q_t grows at the same rate of average trend

²¹Although there is also growth in the q_t series in the years prior to the 1980s, because this growth is driven by other types of equipment than ICT technology, I consider q_t in these prior years to represent a different measure of technological change. As such, I do not use evidence on q_t for these prior years in any of the empirical implementation.

growth observed in 2007-2011 for the next 30 years, and then that growth gradually decays over the subsequent 30 years until completely ceasing in the year 2071, after which q_t is constant.²² With q_t reaching a terminal value, the model will converge asymptotically to a new steady state. Numerically, I solve for the transition path from the initial steady state to this final steady state using a multiple shooting algorithm, where I impose that convergence occurs after a fixed number of periods.²³ As an initial condition, I start the model in the steady state with q_t normalized to 1. Additional details regarding the numerical solution are in Appendix C.

5 Calibration

I now turn my attention to calibration of the model's parameters. Intuitively, the three parameters governing elasticities of substitution between different labor inputs and capital equipment - γ , ρ , and σ - will have crucial implications for the model's predictions regarding the facts described in Section 2.²⁴ However, I momentarily defer a discussion of that calibration and first present the calibration of the other remaining parameters conditional on values for these elasticity parameters.

5.1 Calibration of Other Parameters

Aside from the elasticity of substitution parameters, the parameters to be calibrated are the Cobb-Douglas parameter on structures (α), the share parameters in the production function (μ_j), the depreciation rates for structures and equipment capital (δ_s , δ_e), the discount factor (β), the slope parameter for the Fréchet distribution over occupational skills (θ), the upper bound for the distribution on labor disutility ($\bar{\chi}$), and the functional forms regarding occupation specific productivities ($\phi_j(z)$) and the skill distribution ($F(z)$), with any attendant parameters. Full details regarding the

²²Results are robust, however, to allowing for an even more gradual slowdown in the growth of q_t .

²³In the results presented, I check for convergence 75 years following the end of growth in q_t in 2071. But the results are robust to checking for convergence at a longer time horizon, such as 150 years or more.

²⁴Autor and Dorn (2013) argue how these elasticities of substitution influence the qualitative changes in employment and hourly earnings across different occupations and much of the intuition they develop carries over to this context. As there are some similarities with the capital-skill complementarity framework, the intuition developed in Krusell et al. (2000) is also relevant in this context.

calibration strategy and measurement methods taken to calibrate these are contained in Appendix C; the following reports the general strategy.

The following parameters are all calibrated without numerically solving the model: $\alpha, \delta_e, \delta_s, \beta$. In the model, the share of income paid to structures is given by α , which is thus set to the average share of income in structures in the data for the sample period of 1980-2011. For the depreciation rate on capital equipment, δ_e , I use the value of 0.1 from in Cummins and Violante (2002), and for structures depreciation, δ_s , I use the average depreciation rate implied by the perpetual inventory method applied to the data on structures capital and investment over the sample period 1980-2011. I choose a standard value for the annual discount factor, $\beta = 0.96$.

For the skill distribution and occupation specific productivities, I choose tractable functional forms that are also consistent with the assumptions imposed earlier. Occupational productivity functions take an exponential form, $\phi_j(z) = e^{a_j z}$, consistent with the comparative advantage condition in (4). I assume that general skill follows a normal distribution, with $z \sim N(0, \sigma_z^2)$. The choice of a normal distribution seems to be a reasonable benchmark; the implication of this choice is that the log income distribution will emerge from the combination of a log normal and a Fréchet distribution, which provides a reasonable approximation to the total income distribution.

The remaining parameters are calibrated by solving for the initial steady state and matching certain model to the data in the year 1980. While these parameters are determined simultaneously, I discuss intuition for which moments correspond to which parameters. The share parameters in production, μ_j , are calibrated to match the initial levels of occupational labor shares of income. $\bar{\chi}$ is calibrated to match the value of the employment to population ratio in the year 1980. I calibrate θ to match the cross-sectional variance of log wages in 1980.

To calibrate the occupational productivity parameters, two are calibrated to match the initial hours shares in routine and manual occupations in 1980. For the remaining occupational productivity parameter, there are implicit restrictions on its value to ensure that the comparative advantage condition in (4) is met. To calibrate this remaining parameter, I aim to align the model's earnings across general skill levels with what is observed in the data, as these parameters determine the

relative levels of income across the skill distribution. I do this by generating an empirical wage distribution over skill levels from the data, using the measures of skill as used in constructing Figure 2, and ranking these skill bins by their mean wage in 1980 to obtain a wage distribution over general skill level. In particular, I calibrate the model so that the growth in wages over the general skill distribution matches the growth in wages for the data over two percentile ranges in the middle of the distribution in the year 1980. For this final productivity parameter, I calibrate the model to match the growth in wages from the 10th-25th to the 75th-90th percentiles of the general skill distribution.²⁵ For σ_z^2 , I calibrate the model to match the growth in wages between the 35-50th and 50th-65th percentiles of the general skill distribution. Additional details are available in Appendix C.

5.2 Calibrating Elasticity Parameters

For values of the elasticities of substitution between abstract, routine and manual tasks and equipment, ideally there would be existing (and preferably consensus) estimates in the literature to draw from. Unfortunately, there have not been any attempts yet to estimate these specific parameters. In the absence of such values, as a benchmark, I pursue a calibration approach where I select values for these three elasticity parameters to match changes in hours shares and growth in the capital equipment to hours worked ratio. These changes seem to be reasonable targets, as the elasticity parameters will strongly affect movements in these data moments.²⁶

By calibrating the model to match the hours shares and equipment to hours ratio with only technological change as a driving force, this naturally biases the model's results in favor of technological change, and could be problematic if there are other factors outside the model which are influencing these moments through a different channel. One such prominent factor that could be influencing employment shares and capital equipment growth is the rise of globalization and the

²⁵However, the results are robust to choosing different intervals, or to even considering alternate values of this final productivity parameter within the range of values that is still consistent with the comparative advantage condition.

²⁶Alternatively, these parameters could be calibrated to match movements in the occupational factor shares over this same time period; the results are robust to such a calibration scheme

corresponding increased international trade that has occurred over the past several decades.²⁷ However, as discussed in Katz and Autor (1999) and Autor, Dorn, and Hanson (2013), in the decade of the 1980s, trade and globalization factors appear far less important in shaping labor market outcomes.

I use this time window in the 1980s, when trade forces appear to be less prominent in determining labor market outcomes, and calibrate these elasticity parameters to match changes occurring between 1985 and 1990.²⁸ Specifically, I choose the three elasticity parameters so that the change in the routine hours share, the change in the abstract hours share, and the growth in the ratio of equipment to total hours that occurs along the model’s transition path between the years of 1985 and 1990 exactly match the observed changes in the data for that same time horizon.²⁹ I use 1985 as the starting point for my calibration window to avoid any irregularities stemming from the assumption that the economy is in steady state in 1980.

5.3 Parameter Values

Table 1 lists the calibrated values for all model parameters. The calibration is able to closely match the data and the model moments; results are reported in Appendix C. The values obtained for the three elasticity parameters are $\gamma = 0.2874$, $\rho = -1.0574$, and $\sigma = 0.4632$. Although there are no existing estimates of these elasticities, the estimates of ρ and σ are similar to potentially comparable estimates from Krusell et al. (2000) and Hamermesh (1993). Thus, these parameters obtained through the calibration strategy seem reasonable as a benchmark.³⁰

²⁷This channel is highlighted in Autor, Dorn and Hanson (2013), who show how variation in the exposure to trade shocks from China is correlated with changes in the occupational composition and levels of employment across local labor markets.

²⁸An alternate approach would be to calibrate to match changes across the whole sample period and evaluate the model’s timing of changes relative to the data. Such an exercise reaches the same conclusions as the quantitative analysis presented here.

²⁹In practice, I match the smoothed trends of these statistics to eliminate business cycle frequency fluctuations in calibrating these parameters. Further details for this calibration exercise are given in Appendix C.

³⁰If manual and abstract labor inputs are similar to the skilled and unskilled labor categories used in these studies, then these estimates are directly comparable. These estimates are well within the range of estimates reported in Hamermesh (1993), and somewhat close to the estimates of Krusell et al. (2000), who obtain estimates of $\sigma = 0.401$ and $\rho = -0.495$. The larger negative value for ρ from my calibration may reflect dissimilarities between abstract task input and skilled worker input, the general equilibrium framework used here compared to their partial equilibrium

6 Results

6.1 Hours Shares

Figure 10 presents changes in hours shares in abstract, routine and manual occupations from the data and the model for 1985-2011, where hours shares have been normalized to 0 in the year 1985. Consistent with the calibration, the model matches the trend changes in all three employment shares very well for the years 1985-1990. Further, the model results for the time period of 1990-2000 match quite closely the observed hours share trends in the data. From 2000 on, however, the model's predictions are inconsistent with the observed patterns in the data, with overshooting in the share of hours worked in abstract occupations and undershooting for manual occupations.

6.2 Relative Hourly Earnings

Figure 11 presents the relative average hourly earnings between occupational groups. As there are substantial cyclical fluctuations in the data, I normalize trend values of the model and the data to 1 in 1985.

Both the model and the data show a steady decline in routine hourly earnings relative to abstract hourly earnings, with the model explaining about 50% of the total decline between 1985 and 2011. However, the data series appears to level out in the 2000s, whereas the model continues to show a significant decline. If one compares the decline in relative hourly earnings between 1985 and 2000 or 2005, the model captures about 30-35% of the observed decline in relative hourly earnings. As for hourly earnings in routine occupations relative to manual occupations, the model's results show a near constant ratio of hourly earnings in routine occupations relative to manual occupations, consistent with the patterns in the data.

framework, or the assumption of perfect foresight of the evolution of q_t .

6.3 Employment to Population Ratio

The model's predictions regarding the aggregate employment to population ratio are presented in Figure 12. Consistent with the data, the model observes a fairly stable level for this ratio for the period from 1985-2000. Of particular note is the model's results regarding the decline in employment from 2000 to 2011. That the model is able to account for roughly a fifth of the decline of the employment to population ratio is significant, given that the model neither captures the 2000s recession or the Great Recession.

Figure 13 reports the results for occupational employment rates with the distribution of worker skill in each occupation fixed at 1980, as in Section 2.2.3. As in Figure 12, the model is unable to replicate the full decline in the employment rate for each group, but the results are qualitatively consistent with pattern of employment rate declines increasing as one goes from abstract to routine to manual occupations. However, the magnitudes of the relative declines are substantially different from the data. In particular, the gap between the employment rate for abstract occupation workers and other workers is significantly larger in the model than in the data, especially after 2000.

6.4 Labor's Share of Income

Finally, Figure 14 reports the labor share of income for the model and the data. Similar to the hourly earnings series, as there are substantial cyclical fluctuations in the data, I normalize the data and the model results to their trend values in 1985.

The model predicts about 20% of the total measured decline in the labor share from 1985-2011.³¹ Examining the timing of these declines, however, the model's results regarding the labor share exactly match the data until about the early 2000s. It is only after this point when the model and data diverge, with the model labor share leveling out while it has continued to decline in the data.

Figure 15 presents the decomposition of the labor share into occupational labor shares. The

³¹Notably, this is a smaller final amount than found by Karabarbounis and Neiman (2014), who attribute 50% of the labor share decline to investment specific technological change, though their result is for the global decline of the labor share, and not specifically for the United States.

model's predictions regarding occupational labor shares from 1985 to 2000 are fairly close to the data, and for routine and manual occupations, the model results track the data well through the end of the sample.³² However, after 2000, the model significantly overpredicts the share of income paid to workers in abstract occupations.

6.5 Summary

In summary, the model's results generally compare quite well to the observed data trends through the year 2000. However, it is after the year 2000 that the model's results diverge from the patterns in the data. In particular, the model overstates employment share growth, relative wage growth, and factor income share growth for abstract occupations, as well as overstating employment rate growth for workers with skill levels which originally sorted into abstract occupations. This suggests that technological change seems ill suited to explain this recent decade. Further, these results highlight the potential importance of understanding the observed decline in the demand for abstract tasks in understanding recent declines in aggregate employment and the labor share of income.

Beaudry et al. (2013) also observe a decline in the demand for high skilled or abstract tasks, but suggest that this decline could be accounted for in a similar model framework as the one used here, where there has been a substantial slowdown in technological change. They suggest that the slowing share of nominal GDP attributable to either equipment or ICT investment since the 2000s, shown here in 16, is evidence of such a decline in the rate of technological change. However, my model is able to account for these same patterns in the data. Figure 17 shows that growth in the model's nominal equipment investment share exhibits the same surge in the late 1990s and recent decline through the 2000s as in the data.³³ This decline in the 2000s comes because there is a

³²Notably, the model, while undershooting the share of hours worked in manual occupations, is far more accurate in accounting for the share of income paid to these workers. Some of this discrepancy may be attributable to the composition manual occupations as measured here - see Appendix A for a deeper discussion on the classification and measurement of these occupations.

³³Notably, there appears to be roughly a three year lag between movements in the investment share in the model and in the data. This is a consequence of the perfect foresight approach, as households delay equipment investment to minimize their capital losses on undepreciated equipment. However, the model's primary results are robust to alternative expectations regarding the future time path of q_t , and in the case of completely myopic expectations (believing the present value for q_t is the future value for q_t for all remaining periods), this lag vanishes from the results.

notable slowdown in the growth rate of q_t , which is visible in Figure 18. The key observation here is that although the equipment investment share of nominal GDP is quite flat and declining after the year 2000, the relative *price* of equipment to output is declining at a faster rate, and this is offset by a continued increase in the relative *quantity* of equipment investment and GDP, albeit at a slower pace than before. Thus, although technological change has slowed, it does not appear capable of accounting for this slowing demand for abstract tasks.

7 Underlying Micro Mechanisms

Given the rich heterogeneity at the individual level, it is instructive to understand the underlying changes in individual decisions and how these generate the aggregate patterns shown in Section 6. In discussing these changes, it is important to remember that the model presented here is not a model of worker flows. Thus, the following discussion of changing employment patterns and occupational choice decisions across skill levels should not be interpreted as a discussion of worker switching, but rather in trends over time for all workers of a given skill, encompassing both worker switching decisions and decisions by entrants and exiters from the labor market.

7.1 Occupational Choice and Employment Rates Across Skill Levels

Figure 19 plots the fraction of workers who find it optimal to work in each occupation across general skill level, z , in the initial steady state.³⁴ Consistent with the comparative advantage condition in (4), we see that lower skilled workers tend to choose employment in manual occupations, middle skilled workers tend to choose routine occupations, and higher skilled workers predominantly sort into abstract occupations.

As a result of increased equipment capital accumulation stemming from changes in q_t , relative wages across tasks change over time, inducing workers to change their sorting patterns. Figure 20 shows these changes by plotting the original distribution of skill as in Figure 19, as well as the

³⁴Because of the variation in occupation-specific skills, for each level of general skill, there will be a distribution of occupational choices. See Appendix B for further discussion.

curves dictating the sorting patterns by the end of the sample in the year 2011. By 2011, individuals are far less likely to seek employment in routine task occupations, with lower skilled workers shifting toward manual occupations and higher skilled workers shifting toward abstract occupations. Though these sorting patterns are not directly modeling worker flows, they are consistent with observed worker flow patterns across occupations as described in Cortes (2014) and Groes et al. (2014).

In addition to changes in occupational choice patterns, there are also significant changes in rates of employment over the skill distribution. Figure 21 plots the rates of employment over the skill distribution for the model years 1980 and 2011. The employment rate is upward sloping throughout the skill distribution, with higher skilled individuals being more likely to be employed. Between 1980 and 2011, this employment rate across skill undergoes some “twisting” with employment rates increasing for higher skilled individuals and decreasing for lower skilled individuals. Although wages per efficiency unit are growing slowest for routine occupations, there is not much of a U-shape in employment rate changes across the skill distribution occurring here because of the changing distribution of occupational choice. Namely, the reallocation of middle skilled workers into manual and abstract tasks offsets the slow wage growth for those middle skilled workers still in routine occupations. This pattern of employment rate changes is consistent with evidence in Autor (2010) and Moffitt (2012), who have highlighted that the employment rate declines since 2000 have been most severe amongst lowest skilled workers, not middle-skilled workers.

7.2 Efficiency Units and Wages

In considering how relative hourly earnings across occupations have evolved over time, it is important to think about both changes in the wage paid per efficiency unit of labor and changes in the average efficiency units of labor provided per worker hour in each occupation. The average earnings per hour for occupation j is given by the product of these two terms:

$$Earn_{jt} = w_{jt} \frac{N_{jt}}{H_{jt}}$$

where N_{jt} is the total efficiency units provided in occupation j and H_{jt} is the total hours worked in occupation j .

A relevant question then is how much of the model's changes in relative earnings come from changes in relative wages versus changes in relative efficiency units per hour. Figure 22 plots the model's growth in relative hourly earnings, as in Figure 11, and the growth in just the relative wages per efficiency unit. For both sets of relative earnings, the decline in relative wages is far greater than the decline in relative earnings. This highlights that while the wages of workers in routine occupations are growing slower than in other occupations, the efficiency units per hour of these workers is actually growing fastest.

The potentially surprising result that efficiency units per hour is growing faster in routine occupations than even abstract occupations reflects the ease with which individuals of a given skill can switch occupations. The expansion of hours worked in abstract occupations is primarily coming from individuals who are, on average, less skilled than the existing workers in abstract occupations, as there is no further skill acquisition required for transitioning between occupations. This assumption of no skill acquisition requirements is likely at odds with actual practices in the real world. However, one possible interpretation of the results regarding relative earnings is that these represent the premium for abstract occupational labor apart from the usual skill premium associated with educational attainment or other skill acquisition.

Although there is no change in the distribution over general skill in the model, allowing for such changes is unlikely to change the model's inability to account for the slowdown in demand for abstract occupations after the year 2000. First of all, a natural consideration would be to somehow exogenously account for growth in the supply of college educated persons relative to non-college educated persons. But, as shown in Acemoglu and Autor (2011), this has been very steady since the early 1980s, thus is unclear how this could generate a significant change around the year 2000. Second, if there was some restriction on the ability to provide labor in abstract tasks, potentially through an increased cost of occupational switching or skill acquisition, this would implicitly lead to an increased wage in abstract tasks, but this is at odds with the relative earnings data in Figures 4

and 11.³⁵ Thus, though changes in the distribution of skill over these decades may be important for other outcomes, it is not clear they can account for the slowing hours and earnings paid to abstract occupations.

7.3 Abstract Income Share Decomposition

Given the model's inability to reconcile the observed significant slowdown in the demand for abstract tasks, I conclude by presenting some simple evidence to direct future studies of this slowdown and its causes. In particular, I consider a disaggregation of the abstract task grouping into several subgroups within abstract tasks and examine if this decline in demand is widespread across all abstract occupations, or limited to just a few. I consider 7 occupational subgroups within abstract task occupations; details of this decomposition are available in Appendix A. I measure the slowdown in demand for each of these occupations by examining the annual change in their factor income share, that is, the total income paid to individuals in each occupation divided by total output in the economy.

Figure 23 plots the average annual change in factor income shares for each of these abstract occupation groups, comparing the average annual change between the periods 1980-2000 and 2000-2012. Although the average annual change for abstract occupations has slowed substantially between the periods 1980-2000 and 2000-2012, as seen in Figures 8 and 15, this is not true of all abstract occupations. The most striking changes in demand growth between these two periods are occurring in managerial occupations and high-skilled sales occupations. As a fraction of the total decline in abstract factor income share changes, these two types of occupations contribute 78% of that slowdown. Thus, understanding better the source of the slowdown in demand for abstract tasks seems to be more centered on these groups of abstract occupations.

³⁵Technically, such a restriction in the effective supply of abstract workers could be consistent with these data trends, but only if the efficiency units per hour dropped dramatically since 2000, meaning that the skill level of workers in abstract tasks dropped dramatically in this period. This seems implausible.

8 Conclusion

Much of the discussion surrounding the polarization of the labor market and recent changes in aggregate labor market outcomes has centered around the role of technological change. In this paper, I conduct a formal quantitative evaluation of the role of technological change in explaining these phenomena and find that the model is able to explain a substantial fraction of these changes through the year 2000, but is less able to account for data patterns since then. In particular, the model seems unable to reconcile the slowdown in demand for abstract occupations after 2000, even with a slowdown in the pace of technological change. Further research is needed to determine the exact cause for the apparent slowdown in demand for abstract occupations, though micro level evidence suggests that this slowdown is most prominent amongst managerial and high skilled sales occupations.

References

- ACEMOGLU, D. AND D. H. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” *Handbook of Labor Economics*, 4, 1043–1171.
- AUTOR, D. H. (2010): “The Polarization of Job Opportunities in the US Labor Market: Implications for Employment and Earnings,” *Center for American Progress and The Hamilton Project*.
- AUTOR, D. H. AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103, 1553–1597.
- AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 103, 2121–2168.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): “The Polarization of the US Labor Market,” *American Economic Review: Papers and Proceedings*, 96, 189–194.

- (2008): “Trends in US Wage Inequality: Revising the Revisionists,” *Review of Economics and Statistics*, 90, 300–323.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics*, 118, 1279–1333.
- BEAUDRY, P., D. A. GREEN, AND B. M. SAND (2013): “The Great Reversal in the Demand for Skill and Cognitive Tasks,” NBER Working Paper No. 18901.
- BOSKIN, M. J., E. R. DULBERGER, Z. GRILICHES, R. J. GORDON, AND D. JORGENSEN (1996): *Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee for the Advisory Commission to Study the Consumer Price Index*.
- BURSTEIN, A., E. MORALES, AND J. VOGEL (2014): “Accounting for Changes in Between Group Inequality,” Working Paper.
- CORTES, G. M. (2014): “Where Have the Middle-Wage Workers Gone? A Study of Polarization using Panel Data,” Forthcoming Journal of Labor Economics.
- CORTES, G. M., N. JAIMOVICH, C. J. NEKARDA, AND H. E. SIU (2014): “The Micro and Macro of Job Polarization,” Working Paper.
- CUMMINS, J. G. AND G. L. VIOLANTE (2002): “Investment-specific Technical Change in the United States (1947–2000): Measurement and Macroeconomic Consequences,” *Review of Economic Dynamics*, 5, 243–284.
- DICECIO, R. (2009): “Sticky Wages and Sectoral Labor Comovement,” *Journal of Economic Dynamics and Control*, 33, 538–553.
- EATON, J. AND S. KORTUM (2002): “Technology, Geography, and Trade,” *Econometrica*, 70, 1741–1779.
- ELSBY, M. W., B. HOBIJN, AND A. SAHIN (2013): “The Decline of the US Labor Share,” *Brookings Papers on Economic Activity*.

- FOOTE, C. AND R. W. RYAN (2012): “Labor-market Polarization Over the Business Cycle,” *Federal Reserve Bank of Boston Public Policy Discussion Series*.
- GOOS, M., A. MANNING, AND A. SALOMONS (2014): “Explaining Job Polarization: Routine-biased Technological Change and Offshoring,” *American Economic Review*, 104, 2509–2526.
- GORDON, R. J. (1990): *The Measurement of Durable Goods Prices*, University of Chicago Press.
- (2006): “The Boskin Commission Report: A Retrospective One Decade Later,” *International Productivity Monitor*, 12, 7–22.
- GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (1997): “Long-Run Implications of Investment-Specific Technological Change,” *American Economic Review*, 87, 342–362.
- GROES, F., P. KIRCHER, AND I. MANOVSKII (2014): “The U-Shapes of Occupational Mobility,” Working Paper.
- HAMERMESH, D. S. (1996): *Labor Demand*, Princeton University Press.
- HE, H. AND Z. LIU (2008): “Investment-specific Technological Change, Skill Accumulation, and Wage Inequality,” *Review of Economic Dynamics*, 11, 314–334.
- JAIMOVICH, N. AND H. E. SIU (2012): “The Trend is the Cycle: Job Polarization and Jobless Recoveries,” *NBER Working Paper 18334*.
- JUHN, C. AND S. POTTER (2006): “Changes in Labor Force Participation in the United States,” *Journal of Economic Perspectives*, 20, 27–46.
- JUNG, J. AND J. MERCENIER (2010): “Routinization-biased Technical Change, Globalization and Labor Market Polarization: Does Theory fit the Facts?” *Tepp Working Paper No. 2011-10*.
- KARABARBOUNIS, L. AND B. NEIMAN (2014): “The Global Decline of the Labor Share,” *Quarterly Journal of Economics*, 129, 61–103.


- KATZ, L. F. AND H. DAVID (1999): “Changes in the Wage Structure and Earnings Inequality,” *Handbook of Labor Economics*, 3, 1463–1555.
- KRUSELL, P., L. E. OHANIAN, J.-V. RÍOS-RULL, AND G. L. VIOLANTE (2000): “Capital-skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68, 1029–1053.
- MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (2014): “Has ICT Polarized Skill Demand?: Evidence from Eleven Countries Over 25 Years,” *Review of Economics and Statistics*, 96, 60–77.
- MOFFITT, R. A. (2012): “The US Employment-Population Reversal in the 2000s: Facts and Explanations,” *NBER Working Paper No. 18520*.
- MORIN, M. (2014): “Computer Adoption and the Changing Labor Market,” Working Paper.
- SMITH, C. L. (2013): “The Dynamics of Labor Market Polarization,” Tech. rep., Board of Governors of the Federal Reserve System (US).

Table 1: Calibrated Parameters and Targets

Parameter	Value	Target
σ	0.4632	Growth in capital equipment per hour, 1985-1990
ρ	-1.0574	Change in routine hr. share, 1985-1990
γ	0.2874	Change in manual hr. share, 1985-1990
α	0.1656	Average structures factor share, 1980-2011
μ_m	0.2484	1980 manual labor income share
μ_a	0.4544	1980 abstract labor income share
μ_r	0.5898	1980 routine labor income share
δ_s	0.0330	PIM implied average depreciation, 1980-2011
δ_e	0.1000	Cummins and Violante (2002)
a_m	0.0623	Wage Growth, 10th-25th to 75th-90th percentiles in skill distribution, 1980
a_r	0.1055	Routine Hr. Share, 1980
a_a	0.5314	Abstract Hr. Share, 1980
θ	2.6529	Cross-Sectional Variance in Log Wages, 1980
β	0.96	Standard
σ_z^2	1.7328	Wage Growth, 35th-50th to 50th-65th percentiles in skill distribution, 1980
$\tilde{\chi}$	1.0812	Employment to Population, 1980


Details of measurement for all data targets and the results of calibration target matching are available in Appendix C.

Figure 1: Employment and Wage Polarization Across the Occupational Skill Distribution, 1982-2012




Data comes from the March Annual Supplement to the CPS from 1982-2012 for full-time, full-year wage and salary workers ages 16-64, though results to using all employed over 16. Occupations are ranked by their log mean wage between the years 1982 and 1984, using labor supply weights (the product of weeks and hours worked). Occupations are then evenly divided across 100 percentiles of the wage distribution. The left panel shows employment share growth between 1982 and 2012 and the right panel shows real wage growth, using the PCE price index as a deflator. The smoothed lines come from a locally weighted regression, with bandwidth of 0.8, where the data has been winsorized at a 5% rate. The results, however, are robust to this omitting this procedure in dealing with outliers. The construction of these plots is very similar to those presented in Autor and Dorn (2013).

Figure 2: Log Changes in Employment Rates over the Occupational Skill Distribution from 1982-2012, Worker Skill Distribution in Occupations Fixed at 1980




Data comes from the March Annual Supplement to the CPS from 1982-2012, though results are robust to using an endpoint of 2007. Employment is measured as individuals over the age of 16 who were employed at any point in the previous year. Results, however, are robust to measuring employment as the total number of hours worked in the previous year or even using total employment instead of employment rates. The fraction of workers of a given skill level for each occupation is determined using data on full-time, full-year wage and salary workers age 16-64, though results are robust to including non-full-time, non-full year workers, workers over 64, and the self employed. The occupational skill distribution is identical to the one used in Figure 1. The skill categories used come from five educational categories - less than a HS degree, HS degree, some college, bachelor's degree, and some post bachelor's education - and from nine potential experience categories in 5 year increments - 0-5 years, 5-10 year, and so forth with the top category being 40+ years. Potential experience is calculated as the maximum of age minus years of education minus six and age minus sixteen. The smoothed line shown above comes from a locally weighted regression, with bandwidth of 0.8, where the data has been winsorized at a 5% rate. The results, however, are robust to this omitting this procedure in dealing with outliers.

Figure 3: Hours Shares of Abstract, Routine, and Manual Occupations




Data comes from the March Annual Supplement to the CPS from 1975-2012 on all employed persons (including self-employed) in the non-farm business sector, though similar patterns are observed if one restricts the sample to be 16-64 full time, full-year wage and salary workers. Occupational classifications come from Autor and Dorn (2013). Additional details regarding data construction and sources are available in Appendix A.

Figure 4: Relative Average Hourly Earnings Across Occupational Groups




Data is from the CPS March Supplement files on employed persons in the non-farm business sector. Occupational classifications come from Autor and Dorn (2013). Average hourly earnings is constructed by dividing total wage and salary income for an occupational group by total hours worked by individuals in that occupation. Relative hourly earnings is normalized to 1 in 1975. Additional details regarding data construction and sources are available in Appendix A.

Figure 5: Employment to Population Ratio, 1975-2012




Data comes from the March Annual Supplement to the CPS from 1975-2012. Employment is measured as in Figure (2). Additional details regarding data construction and sources are available in Appendix A.

Figure 6: Labor's Share of Income, 1975-2013




Data for the labor share of income for the non-farm business sector comes from the BLS Labor Productivity and Costs series.

Figure 7: Change in Employment Rate by Occupation 1975-2012, Worker Skill Distribution in Occupations Fixed at 1980




Data comes from the March Annual Supplement to the CPS from 1975-2012. Employment rates as measured as in Figure (2). The employment rate for each occupation has been normalized to 1 in the year 1975 to facilitate comparison between occupations.

Figure 8: Occupational Shares of Income, 1975-2012




Occupational income data comes from the March CPS Supplement and data on nominal GDP comes from the BLS Labor Productivity and Costs series. Data is for the non-farm business sector. The occupational share of income for occupation j is defined as total wage and salary income paid to a worker in occupation j divided by nominal GDP. As such, the sum of all occupational shares of income equals the aggregate labor share of income. Additional details regarding data sources and construction are available in Appendix A.

Figure 9: Time Path for q_t




The series for q_t is generated using quality adjusted equipment prices from DiCecio (2009) and quality adjusted consumption prices for non-housing, non-energy services and non-durable consumption. Unadjusted consumption prices are available from the BEA; the adjustment for quality bias is detailed in Appendix A. The series is smoothed with an HP filter with parameter 6.25.

Figure 10: Hours Shares 1985 to 2011, Model vs. Data




The hours share for routine occupations is in red, manual in green, and abstract in blue. Model outcomes are denoted with dashed lines. All series are normalized to 0 in 1985.

Figure 11: Relative Hourly Earnings Between Occupations 1985 to 2011, Model vs. Data




Average hourly earnings in routine occupations divided by hourly earnings in abstract occupations is shown in blue; hourly earnings in routine occupations divided by hourly earnings in manual occupations is shown in green. Model results are in dashed lines. The series are normalized by their trend value in 1985, where the trend is obtained with an HP filter with parameter 6.25.

Figure 12: Employment to Population Ratio, Model vs. Data, 1985 to 2011




Data on the employment to population ratio is the same as in Figure 5. The employment to population ratio for the model is simply total employment, as the population is constant at 1.

Figure 13: Occupational Employment Rates, Skill to Task Mapping Fixed at 1980 Level




Occupational employment rates in the model are constructed as in the data - the weighted average of employment rates of skill types employed in each occupation.

Figure 14: Labor's Share of Income 1985 to 2011, Model vs. Data




Data on the labor share is simply the annual average of the series used in Figure 9; all series are normalized by subtracting off their trend value at 1985, where the trend is computed with an HP filter with parameter 6.25.

Figure 15: Occupational Labor Shares of Total Income 1985 to 2011, Model vs. Data




Occupational labor shares have been constructed as in Figure 8. All series are normalized to 0 in the year 1985 by subtracting off their trend value at 1985, where the trend is computed with an HP filter with parameter 6.25.

Figure 16: Equipment Investment and ICT Investment as a Share of Nominal GDP, 1947-2014




Data on nominal equipment investment comes from adding investment in equipment from NIPA Table 1.1.5 to software investment from NIPA Table 5.3.5. Data on nominal ICT investment comes from adding information processing equipment investment and software investment from NIPA Table 5.3.5. Data on nominal GDP is from NIPA Table 1.1.5.

Figure 17: Growth in the Equipment Investment Share of Nominal GDP: Model vs. Data




Data on the equipment investment share of nominal GDP is the annualized version of the data used in Figure 16. The model equipment investment share of nominal GDP is given by $\frac{\frac{1}{q_t}(K_{et+1} - (1 - \delta_e)K_t)}{Y_t}$. Both series are normalized to 1 in 1985 by dividing by the initial investment share in 1985 to facilitate comparison.

Figure 18: Annual Growth Rate of q_t




Growth in q_t is computed by $\log(q_{t+1}) - \log(q_t)$.

Figure 19: Occupational Choice Distribution Across General Skill, 1980




The share of workers of a given skill level z who choose to be employed in occupation j is given by: $\frac{(w_{jt}e^{a_j z})^\theta}{\sum_i (w_{it}e^{a_i z})^\theta}$. Additional detail is given in Appendix B. The green shaded region denotes workers in manual occupations, the red denotes routine, and the blue denotes abstract.

Figure 20: Occupational Choice Distribution Across Skill, 1980 and 2011




The share of workers of a given skill level z who choose to be employed in occupation j is given by: $\frac{(w_{jt} e^{a_j z})^\theta}{\sum_i (w_{it} e^{a_i z})^\theta}$. Additional detail is given in Appendix B. The green shaded region denotes workers in manual occupations in 1980, the red denotes routine, and the blue denotes abstract. The light green dashed line indicates the share of workers employed in manual task occupations in the year 2011, and the light blue dashed line indicates the share of workers employed in either routine or manual task occupations in the year 2011.

Figure 21: Rates of Employment Across the Skill Distribution, 1980 and 2011




The employment rate for a given skill level z' is given by $\int_{\omega} \min\left(\frac{\omega_t}{C_t} / \bar{\chi}, 1\right) d\Phi_t(\omega_t | z)$, where $\Phi_t(\omega_t | z)$ is the distribution function for a random variable distributed Fréchet $(-\sum_j (w_j \phi_j(z))^{-\theta}, \theta)$. This is discussed in greater detail in Appendix B. The solid line shows the employment rate across skill for the year 1980; the dashed line shows it for the year 2011.

Figure 22: Changes in Relative Earnings Per Worker vs. Changes in Relative Wages Per Efficiency Unit



The blue lines show earnings per hour and wages per efficiency unit in routine occupations divided by those in abstract occupations. The green line show earnings per hour and wages per efficiency unit in routine occupations divided by manual occupations. Earnings per hour in occupation j are given by $\frac{w_{jt}N_{jt}}{H_{jt}}$ and relative earnings are denoted with dashed lines. Wages per efficiency unit in occupation j are simply w_{jt} and relative wages are shown with solid lines.

Figure 23: Average Annual Change in Factor Income Share for Abstract Occupations, 1980-2000 vs. 2000-2012



The factor income share for an occupation j is given by the total income paid to individuals in that occupation divided by GDP, $\frac{\text{Total Income}_j}{\text{GDP}}$. White bars show the average change in this share from 1980-2000, and black bars show the average change in this share from 2000-2012. Individual abstract occupations are generated from occupational classifications as described in Appendix A. By construction, the sum of average changes across all these occupations equals the total average change for all abstract occupations over these time windows. Further, although these results intersect directly the late 1990s and dot com bubble, as well as the Great Recession, the qualitative conclusions of these results are robust to omitting these time windows.

Appendix A: Data Sources and Construction

For data on changes in the distribution of occupational employment and hourly earnings, I use data from the Annual March CPS supplement files. For the construction of the relative price of equipment investment to consumption, I use detailed price series on consumption and investment from the National Income and Product Accounts. Before describing the procedures for preparing the data, however, I consider the occupational code system obtained from Autor and Dorn (2013).

Occupational Classifications

For time-consistent occupational classifications, I use the modified occupational classification system used in Autor and Dorn (2013). However, the occupational classification system does not account for the recent 2010 revision to the census occupational codes. Thus, I use crosswalks from the census bureau to extend this classification system forward as needed. A full concordance of this extension method is available upon request.

I make one other modification to the above mentioned classification system. While this system generally does an excellent job of providing time consistent measures of occupations, there are still subtle discontinuities each time the classification systems are revised. The problem with these discontinuities is determining what the appropriate growth in employment would have been absent the discontinuity. To determine the growth across these periods, I utilize the timing of the CPS survey, which has individuals both report their present occupation and their past year's occupation. While I primarily use the past year's employment to be aligned with the data on income, I use the present year data to determine the growth in employment across time periods with a classification discontinuity, as the discontinuity affects the present year series a year later. With the growth in employment across the discontinuity period, I project backward one period what employment growth should have been based on the prior survey's present year data, and then scale up past years to be consistent with change in classification.

To map occupations into the three task classes - abstract, routine and manual - I use the mapping

from Autor and Dorn (2013) which looks at task intensity across occupations and determines which occupations correspond to which classes based on whether or not the task intensity scores are above average. With this mapping, abstract, routine and manual occupations are identified as follows:

Abstract: management, professional, technical, financial sales and public safety occupations

Routine: precision production and craft occupations, machine operators, assemblers and inspectors occupations, and clerical/administrative support and retail sales occupations.

Manual: transportation, construction, mechanics, mining occupations and low skill service occupations

It is worthwhile to note that there have been a wide variety of mappings employed to relate occupations to tasks. Alternate classification systems have often additionally distinguished jobs by two categories: routine vs. non-routine and cognitive vs. manual.³⁶ As in Autor and Dorn (2013), abstract occupations are essentially non-routine cognitive jobs and manual jobs are non-routine manual jobs, where as routine jobs are either routine cognitive or routine manual. Per the occupational classifications of Autor and Dorn (2013), occupations defined as manual occupations are divisible into two groups - low skill service occupations and construction, mechanic, repairer, mining and transportation occupations. These latter manual occupations are sometimes classified as routine occupations, as they have both strong manual and routine content. Further, the labor market outcomes for this subset of manual occupations have not always been similar to other manual occupations, particularly with the construction boom preceding the 2007-2009 recession. That said, I follow Autor and Dorn (2013) and label these as manual jobs, although Jaimovich and Siu (2012) consider these to be routine occupations as well, and Smith (2013) labels these occupations as “other.” However, using a classification scheme as in Jaimovich and Siu (2012) generates broadly comparable results.

For the decomposition of abstract occupations in Section 8, I use the following decomposition of abstract occupations across the occupational coding system described in Dorn (2009):

³⁶Acemoglu and Autor (2011) use this particular classification scheme.

Abstract Occupation Group	Examples	Codes (Dorn [2009])
Managerial and Mgmt. related	Managers, Accountants, Mgmt Analysts	4-37
Engineering and Technician	Engineers, Architects, Drafters, Lab techs., Pilots	43-59, 203-235
Computer and Physical Sciences	Computer Scientists, Chemists, Physicists	64-76
Biomedical	Doctors, Dentists, Veterinarians, RNs, Therapists	77-106
Education, Arts and Entertainment	Teachers, Writers, Designers, Actors, Reporters	183-199
Social Sci., Law and Social Serv.	Economists, Psychologists, Soc. Workers, Lawyers	166-178,417-423
High Skilled Sales	Sales Proprietors, Fin./Ins. Sales, Real Estate Agents	243-258

CPS March Supplement Files

For hourly earnings data, I use the Annual March CPS Supplement files from 1976 to 2012, obtained through IPUMS (King et al. (2010)), which reports information on individuals' income for the prior year. I restrict my attention to workers employed in the non-farm business sector, and thus drop workers if they are employed outside of this sector. My focus on average hourly earnings is generally consistent with the recent literature.³⁷ It is also consistent with constructing aggregate factor shares for occupational labor income.

The March CPS files require several adjustments for hours and income responses. For hours non-respondents, I impute hours worked as in Krusell et al. (2000). For income, there are a variety of potential measurement issues, the most prominent being the top-coding that is applied in the data. Income recorded in the CPS is subject to two levels of top-coding - public use data top-coding and internal top-coding. In the data prior to the survey year 1996, income levels exceeding the public top code are simply listed at the top code value. From 1996 onwards, income levels at the top code are replaced with an average of actual income values for individuals with similar characteristics. Larrimore et al. (2008) provide comparable measures for this "cell-mean" procedure dating back to the 1976 survey, so I use these cell-means until the CPS begins using them in practice. However, this procedure does not provide any adjustment to compensate for the internal

³⁷See Footnote 8 in Acemoglu and Autor (2011) for a more in-depth discussion of using hourly earnings or average wages.

top-coding. To make this correction, I fit a Pareto distribution to the tail of the income distribution, using the 90th and 95th percentiles, and use it to impute a correction term to the cell-means to compensate for lumpy changes induced by the internal top code.

To obtain a comparable aggregate measure of the labor share from the household survey, I need to impute earnings of the self-employed. I follow the BLS procedure of imputing self-employed wage earnings as the average wage (over all observations where the implied wage is greater than half the minimum wage) multiplied by the usual number of hours worked. In cases where individuals both both business income and wage income, I imputed self-employment earnings as the average wage times either $2/3$ the total hours worked if the primary job was self-employment, and $1/3$ the total hours worked if the primary job was wage and salary employment. All results are robust to varied modifications in this imputation procedure.

Even with the imputed self-employed income, the aggregate income series derived from the March supplement files will be unlikely to completely match the aggregate income series the BLS calculates from the establishment survey. This is because income in the establishment survey also includes employer contributions to health care, as well as payments in kind (particularly, the exercise of non-qualified stock options). As a result, the estimated aggregate income series from the household is always short about 80-85% of the total income reported by the BLS; however, despite the level gap, the dynamics of the two series match very closely. As such, to match the aggregate labor share, I scale up the income series for each occupational group to match the aggregate total. Given that the limited data from the CPS regarding employer contributions to health care shows similar patterns across occupations as do wage and salary incomes, this assumption seems quite reasonable.

The final adjustment that is made, though does ultimately not affect the final results, is to make a correction for the changing in the weighting scheme implemented in 2003. Aggregate statistics reported by the BLS adjust the population weights to shift this discontinuity back in time to 2000; I employ a comparable procedure to redistribute population weights consistently between 2000-2003.

With these adjustments, I am able to construct hours worked by occupation and average hourly earnings by occupation. I compute the aggregate employment to population ratio as simply the total of employed persons over age 16 divided by the civilian population over age 16. To compute occupational labor shares, I take the total income in each occupation, divide it by total nominal output for the non-farm business sector, as measured in the BLS Labor Productivity and Costs Series, and adjust the shares so that the total labor share matches the aggregate measure, as mentioned above.

Relative Price Data

I use the quality adjusted price series for the price of capital equipment investment, the same as used in DiCecio (2009). For data on the price of consumption, I use price indices for Personal Consumption Expenditures and construct a Tornqvist index for the price of non-durable consumption and non-housing, non-energy services, as in Cummins and Violante (2002). I also go one step further and correct for price measurement bias in the consumption prices using the recommendations of Boskin et al. (1996) and Gordon (2006). These corrections require delicacy to apply, as the original findings on consumption price quality bias are for the CPI and not the PCE. But a large fraction of the underlying data in the PCE comes from the CPI and by mapping different consumption subcategories between the underlying data for each of these, I can adjust the growth rate of prices in the consumption subcategories where quality bias is identified. Further details of this procedure and the exact adjustments made are available upon request. However, all results in the paper are robust to using the BEA series without bias adjustment.

Appendix B: Model Equilibrium Conditions

Here I report the full set of equilibrium conditions for the model, however, prior to reporting these, it is helpful to present several of the household's decision rules and outcomes.

First, given the separability of preferences over consumption and labor supply at the individual

level, and the utilitarian maximization of household utility, the household will equalize consumption across individuals. That is, $C_t(s) = C_t$.

Second, recall that potential income, the income a worker obtains if employed, is given by $W_t(s) = \max_j w_{jt} z^j \phi_j(z)$. The household chooses to have an individual work if the marginal gain in household utility from employment of that individual exceeds the marginal cost of that individual's labor disutility. Thus, an individual i works only if $\frac{W_t(s_i)}{C_t} > \chi_i$. Given that the disutility of work follows a uniform distribution, given an individual's general and occupational skills, the probability that said worker i is employed is given by $P(\chi_i < \frac{W_t(s_i)}{C_t})$. With the independence assumption on the worker skills and disutility, this is simply given by $\min\left(\frac{W_t(s_i)}{C_t}/\bar{\chi}, 1\right)$.

Further, it is helpful to think about both the distribution of the employment and occupational choice decisions conditioning only on the general skill, z . Because individual occupational skills follow a Fréchet distribution, potential income conditional on z , $W_t | z$ will also follow a Fréchet distribution, with $W_t | z \sim \text{Fréchet}(\sum_j (w_{jt} \phi_j(z))^\theta, \theta)$.³⁸ This result simply follows from the property that the maximal order statistic of independent Fréchet random variables is also distributed Fréchet.³⁹ Using $\Phi(W_t | z)$ to denote the distribution function for potential income conditional on z at time t , the employment rate for individuals conditional on skill z is given by:

$$\int_0^\infty \min\left(\frac{W_t}{C_t}/\bar{\chi}, 1\right) d\Phi(W_t | z)$$

which is simply the average employment rate across all individuals with general skill z . Thus, for an individual with skill z , changes in the employment rate over time will driven that individ-

³⁸The notation for the Fréchet distribution parameters is the same as used in Eaton and Kortum (2002), with $-\sum_j (w_j \phi_j(z))^{-\theta}$ being the location parameter and θ being the slope parameter.

³⁹Potential income, conditional on z , is simply the maximal order statistic of the three potential incomes in each distribution. Thus, the distribution function for $W(z)$ is given by $P(W(z) \leq \omega) = P(w_{jt} z^j \phi_j(z) \leq \omega \quad \forall j | z)$. Because the individual occupational skill values are independent of each other, this is simply:

$$\Pi_j P(w_{jt} z^j \phi_j(z) \leq \omega | z) = \Pi_j P\left(z^j \leq \frac{\omega}{w_{jt} \phi_j(z)} | z\right) = \Pi_j e^{-\left(\frac{\omega}{w_{jt} \phi_j(z)}\right)^{-\theta}} = e^{-\sum_j (w_j \phi_j(z))^\theta \omega^{-\theta}}$$

which is simply the distribution function for a random variable distributed Fréchet with location parameter $-\sum_j (w_j \phi_j(z))^{-\theta}$ and scale parameter θ .

ual's marginal contribution to household disutility (driven primarily by changes in the wage) and changes in distribution for disutility of work. The aggregate employment rate can then simply be determined by averaging up over all skill levels z .

By this same reasoning, analyzing the occupational choice decision, it is not hard to show that for an individual with skill level z , the probability of choosing to work in occupation j at time t is given by⁴⁰:

$$\frac{(w_{jt}\phi_j(z))^\theta}{\sum_i (w_{it}\phi_i(z))^\theta}$$

With these rules from the household, it is straightforward to list the equilibrium conditions for the model. First, there is the production function and the first order conditions from the firm's problem:

$$Y_t = K_{st}^\alpha \left[\mu_m N_{mt}^\sigma + (1 - \mu_m) \left[(1 - \mu_a) \left[(1 - \mu_r) K_{et}^\gamma + \mu_r N_{rt}^\gamma \right]^{\frac{p}{\gamma}} + \mu_a N_{at}^p \right]^{\frac{\sigma}{p}} \right]^{\frac{(1-\alpha)}{\sigma}}$$

$$w_{mt} = (1 - \alpha) \mu_m K_{st}^\alpha \left(\frac{Y_t}{K_{st}^\alpha} \right)^{1 - \frac{\sigma}{1-\alpha}} N_{mt}^{\sigma-1}$$

$$w_{rt} = (1 - \alpha) (1 - \mu_m) (1 - \mu_a) \mu_r K_{st}^\alpha \left(\frac{Y_t}{K_{st}^\alpha} \right)^{1 - \frac{\sigma}{1-\alpha}} N_{rt}^{\gamma-1} \Omega_t^{\frac{\sigma}{p}-1} \Theta_t^{\frac{p}{\gamma}-1}$$

$$w_{at} = (1 - \alpha) (1 - \mu_m) \mu_a K_{st}^\alpha \left(\frac{Y_t}{K_{st}^\alpha} \right)^{1 - \frac{\sigma}{1-\alpha}} N_{at}^{p-1} \Omega_t^{\frac{\sigma}{p}-1}$$

⁴⁰The probability that an individual chooses occupation j is given by the probability that potential income in occupation j exceeds the maximum of potential income in the other two occupations. From this perspective, it is not hard to show that the result above. Eaton and Kortum (2002) prove a similar result in a different context, where these properties of the Fréchet distribution are used to identify the probability that a country provides a given good at a lower price than all other countries.

$$r_{st} = \alpha \frac{Y_t}{K_{st}}$$

$$r_{et} = (1 - \alpha)(1 - \mu_m)(1 - \mu_a)(1 - \mu_r) K_{st}^\alpha \left(\frac{Y_t}{K_{st}^\alpha} \right)^{1 - \frac{\sigma}{1 - \alpha}} K_{et}^{\gamma - 1} \Omega_t^{\frac{\sigma}{p} - 1} \Theta_t^{\frac{p}{\gamma} - 1}$$

where Ω_t and Θ_t are CES nests from the production function:

$$\Omega_t = (1 - \mu_a) \left((1 - \mu_r) K_{et}^\gamma + \mu_r N_{rt}^\gamma \right)^{\frac{p}{\gamma}} + \mu_a N_{at}^p$$

$$\Theta_t = (1 - \mu_r) K_{et}^\gamma + \mu_r N_{rt}^\gamma$$

From the household side, there are two Euler equations, one for each type of capital:

$$\frac{1}{C_t} = \beta \frac{1}{C_{t+1}} (1 - \delta_s + r_{st+1})$$

$$\frac{1}{q_t C_t} = \beta \frac{1}{C_{t+1}} \left(\frac{1 - \delta_e}{q_{t+1}} + r_{et+1} \right)$$

And finally, there are the labor market clearing conditions (expressed here in terms of income) and the aggregate resource constraint:

$$w_{jt} N_{jt} = \int_z \int_0^\infty \frac{(w_{jt} \phi_j(z))^\theta}{\sum_i (w_{it} \phi_i(z))^\theta} W_t \min \left(\frac{W_t}{C_t} / \bar{\chi}, 1 \right) d\Phi(W_t | z) dF(z) \quad \text{for } j \in \{a, r, m\}$$

$$C_t + K_{st+1} + (1 - \delta_s) K_{st} + \frac{1}{q_t} (K_{et+1} - (1 - \delta_e) K_{et}) = r_{et} K_{et} + r_{st} K_{st} + \int_z \int_0^\infty W_t \min \left(\frac{W_t}{C_t} / \bar{\chi}, 1 \right) d\Phi(W_t | z) dF(z)$$

Appendix C: Solution and Calibration Methods

Solution Method

I use a multiple shooting algorithm to solve for the transition path between two steady states of the model economy. The exact solution algorithm is as follows:

1. Given values for the model parameters, solve for the initial and terminal steady states.
2. Use the initial steady state value of consumption as a guess for the value of consumption in the first period and set upper and lower bounds around this guess (these bounds are checked after the fact to ensure they are not hit).
3. Given the initial guess for consumption and capital stocks from the initial steady state, solve the model for period 1 and then iteratively solve the model forward along the entire time path of shocks and then an extended period of N periods where the shocks are held constant at their terminal values.
4. Check the value of consumption at the end of the path and the extended N periods and compare it to the terminal steady state value of consumption. If consumption is too high, update the guess by choosing the midpoint of the lower bound and the current guess; similarly if consumption is too low (standard monotonicity argument). Employ this bisection algorithm until the difference in guesses in adjacent periods has converged.
5. Now given a solution for the first period's consumption, guess a value for the second period's consumption given the capital investment decisions from period 1, and solve the time path again. Obtain a value for consumption and capital investment decisions in period 2.
6. Perform this solution routine to obtain a time series for consumption and capital stocks for the requisite number of periods.

In theory, a simple single shooting algorithm could work, and the iterative procedure would not be needed, but as is well known, single shooting is known to be unstable in explosive saddle-

point systems and accumulating numerical errors may creep in with successive iterations that make convergence near impossible (see Lipton et al. (1983) for a discussion of this and multiple shooting in economics). In addition to stability, the further advantage of an iterative procedure such as this is manifest in calibration when the targets in calibration are changes along the transition path. To calibrate the model, only the number of periods needed for the calibration need to be solved, and not the entire time path, thus reducing computational time.

Additionally, solving the model requires evaluating multidimensional integrals without a given closed form solution. As a result, I use a Gauss-Legendre quadrature to evaluate these integrals.

Calibration Approach

Steady State

The values for the Cobb-Douglas share of structures (α) and the depreciation rates (δ_s, δ_e) are obtained without numerically solving the initial steady state.

To calibrate α , I use the fact that along a perfect foresight path, ex post returns to capital structures and capital equipment are equated as follows:

$$(1 - \delta_s) + r_{st+1} = \frac{q_t}{q_{t+1}} (1 - \delta_e) + q_t r_{et+1}$$

Additionally, the constant returns to scale nature of the production function implies:

$$\frac{r_{st} K_{st} + r_{et} K_{et}}{Y_t} = 1 - Lsh_t$$

where Lsh_t is the total labor share of income at date t . Given data on the labor share, nominal output, the time path of q_t , and values for the depreciation rates, these equations yield two equations in two unknowns and can be solved to find time paths of the rental rates in the data. I then construct α as:

$$\alpha = \overline{\left(\frac{r_{st} K_{st}}{Y_t} \right)}$$

where the average is taken over the period 1980 to 2011.

Data for capital equipment and capital structures comes from the BEA data on Fixed Assets for the non-farm private sector. For structures capital, real quantities have been constructed using BEA price deflators; for equipment, I follow the procedure set forth in Cummins and Violante (2002) and use the perpetual inventory method with the quality adjusted prices to obtain a measure of capital equipment.

To calibrate the depreciation rate for structures, I back out the depreciation rate from the capital accumulation equation, and take the average over the period 1980-2011:

$$\delta_s = 1 - \overline{\left(\frac{K_{st+1} - I_{st}}{K_{st}} \right)}$$

The remaining parameters to calibrate by numerically solving the initial steady state are: the share parameters in the CES nests (μ_j), the support for the distribution of the disutility of labor ($\bar{\chi}$), the productivity parameters in occupational specific productivities (a_j), the Fréchet parameter on occupation specific skill (θ), and the variance of the skill distribution (σ_z^2). These parameters are calibrated simultaneously to match a series of targets, but I give some intuition as to which moments correspond to which parameters. The μ_j parameters are calibrated to match the initial factor shares for different occupational types of labor in 1980. $\bar{\chi}$ is calibrated to match the initial employment to population ratio in 1980. θ is calibrated to match the cross-sectional variance of log wages in 1980.

Two of the a_j parameters are calibrated to match the initial hours distribution in 1980, and the final parameter for occupational productivity, as well as the variance of the general skill distribution, are calibrated to match a target on wage growth over the general skill distribution in 1980. To construct a wage distribution over general skills in the data, I rank skill populations according to

their mean log wage based, where skill populations are defined over nine potential experience and five educational groups. I then develop a one hundred percentile distribution by evenly dividing up each skill group based on its total weight in the overall working population. In short, the procedure for generating said wage distribution over skill is near identical to the one used to generate the occupational skill distribution for Figure 1. I then look at average log wage in the 10-25th, 35-50th, 50-65th, and 75-90th percentiles and then compute the difference across the 10-25th and 75-90th and the 35-50th and 50-65th percentiles, and calibrate the model to match these targets, using the model's general skill distribution.

Along the Transition Path

There are three parameters to be calibrated along the transition path - the three elasticity of substitution parameters (γ, ρ, σ). To calibrate along the transition path, I simply solve for the values for the hours shares, the employment to population ratio and capital equipment to hours from the steady state until the year 1995.⁴¹ I smooth the targets and the model moments using an HP filter with parameter 6.25 to remove business cycle frequencies from 1980-1995 and thus calibrate to match the change in the smoothed series between 1985 and 1990.

⁴¹Total hours for the equipment to hours ratio are measured from the CPS as total hours worked by individuals over age 16.

Calibration Results

Moment	Data	Model
Smoothed growth in capital equipment per hour, 1985-1990	1.3345	1.3345
Smoothed change in routine hr. share, 1985-1990	-0.0184	-0.0184
Smoothed change in abstract hr. share, 1985-1990	0.0218	0.0218
1980 manual labor income share	0.1628	0.1626
1980 abstract labor income share	0.2493	0.2542
1980 routine labor income share	0.2232	0.2225
Routine Hr. Share, 1980	0.4012	0.4068
Abstract Hr. Share, 1980	0.2955	0.2854
Employment to Population, 1980	0.6846	0.6844
Cross-Sectional Variance in Log Wages, 1980	0.3653	0.3653
Wage Growth, 35th-50th to 50th-65th percentiles in skill distribution, 1980	0.0770	0.0951
Wage Growth, 10th-25th to 75th-90th percentiles in skill distribution, 1980	0.5319	0.5273