Journal of Monetary Economics 91 (2017) 69–87

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmoneco

Disappearing routine jobs: Who, how, and why?[☆]

Guido Matias Cortes[a,b], Nir Jaimovich[c,d,*], Henry E. Siu[d,e]

[a] Department of Economics, York University, 4700 Keele Street, Toronto, ON M3J 1P3, Canada
[b] Rimini Centre for Economic Analysis, 75 University Ave. W., Waterloo, ON N2L 3C5, Canada
[c] University of Zurich, Department of Economics, Schönberggasse 1, CH-8001 Zürich, Switzerland
[d] National Bureau of Economic Research, 1050 Massachusetts Ave., Cambridge, MA 02138, USA
[e] Vancouver School of Economics, University of British Columbia, 6000 Iona Drive, Vancouver, BC V6T1L4, Canada

ARTICLE INFO

Article history:
Received 29 November 2016
Revised 6 September 2017
Accepted 6 September 2017
Available online 14 September 2017

Keywords:
Routine occupations
Job polarization
Automation
Labor force participation

ABSTRACT

We study the deterioration of employment in middle-wage, routine occupations in the United States in the last 35 years. The decline is primarily driven by changes in the propensity to work in routine jobs for individuals from a small set of demographic groups. These same groups account for a substantial fraction of both the increase in non-employment and employment in low-wage, non-routine manual occupations observed during the same period. We analyze a general neoclassical model of the labor market featuring endogenous participation and occupation choice. In response to an increase in automation technology, the framework embodies a tradeoff between reallocating employment across occupations and reallocation of workers towards non-employment. Quantitatively, we find that this standard model accounts for a relatively small portion of the joint decline in routine employment and associated rise in non-routine manual employment and non-employment.

1. Introduction

In the past thirty-five years, the US economy has seen a sharp drop in the fraction of the population employed in middle-skilled occupations. This employment loss is linked to the disappearance of *routine* occupations—those focused on a relatively narrow set of job tasks that can be performed by following well-defined instructions and procedures. This fall in per capita routine employment is a principal factor in the increasing polarization of the labor market, as employment shares have shifted toward the top and bottom tails of the occupational wage distribution. Autor et al. (2003) and the subsequent literature suggest that job polarization is due to progress in automation technologies that substitute for labor in routine tasks.

In spite of the large and growing literature on polarization, relatively little is known regarding the process by which routine occupations have declined. This is true with respect to who the loss of routine job opportunities is affecting most acutely, and how they have adjusted in terms of employment and occupational outcomes.[1] And though the number of stud-

☆ We thank David Dorn, Richard Rogerson, and numerous seminar and workshop participants for helpful comments. Siu thanks the Social Sciences and Humanities Research Council of Canada for support.

* Corresponding author at: University of Zurich, Switzerland.
E-mail addresses: gmcortes@yorku.ca (G.M. Cortes), nir.jaimovich@uzh.ch (N. Jaimovich), hankman@mail.ubc.ca (H.E. Siu).

[1] An important exception is Autor and Dorn (2009) who consider changes in the age composition of different occupations. Cortes (2016) analyzes transition patterns out of routine occupations and the associated wage changes experienced by workers.

http://dx.doi.org/10.1016/j.jmoneco.2017.09.006
0304-3932/

ies is growing, the quantitative role of progress in automation technology in the aggregate decline of routine employment is also unresolved. This paper contributes to these questions.

In Section 2, we study the proximate empirical causes of the decline in per capita employment in routine occupations. In an accounting sense, roughly one-third of the fall observed in the past four decades is due to demographic compositional change within the US population. The more important factor is a sharp change in the propensity of individuals of given demographic characteristics to work in routine jobs. These composition and propensity effects are strongest for a relatively small set of demographic groups. As a result, the vast majority of the fall in routine employment can be accounted for by changes experienced by individuals of specific demographic characteristics.

For routine manual occupations, this is the group of young and prime-aged men with low levels of education. Increasing educational attainment and population aging in the US means that the fraction of individuals with these characteristics is falling. Moreover, these same demographic groups have experienced the sharpest drops in the propensity for routine manual employment. For routine cognitive occupations, the vast majority of the decline is accounted for by changes in employment propensities of young and prime-aged women with intermediate levels of education.

We next document the labor market outcomes that offset this fall in per capita routine employment. For the key demographic groups identified above, we see an increase in the propensity for non-employment (unemployment and labor force non-participation) and employment in non-routine manual occupations. These changes are relevant for two recent phenomena: (i) rising non-employment, and (ii) the reallocation of labor to low-wage occupations within the U.S. working-aged population. We show that the propensity changes of the key demographic groups responsible for the decline of routine employment also account for large fractions of the changes in (i) and (ii).[2]

In the remainder of the paper, we explore the role of advances in automation technology in accounting for these phenomena. We do so within the context of a general, flexible neoclassical model of the labor market featuring endogenous participation and occupational choice, presented in Section 3. Our main findings can be summarized as follows. In Section 4 we demonstrate analytically that advances in automation cause workers to leave routine occupations and sort into non-employment and non-routine manual jobs. We then show that the neoclassical framework embodies an important tradeoff: generating a role for increased automation in reallocating employment from routine to non-routine manual occupations comes at the expense of automation's role in reallocation from employment to non-employment, and vice versa.

Section 5 discusses the quantitative specification of the model. In Section 6, we find that advances in automation technology—as measured by the increase in the stock of ICT (information and communications technology) capital valued in efficiency units—on its own accounts for a relatively small portion of the joint decline in routine employment and associated rise in non-routine manual employment and non-employment. While this result relates to the quantitative exercise we consider, it is worth stressing that the theoretical approach that we outline in Sections 3 and 4 is a useful, general framework for future investigation.

With regards to the literature, our theoretical work is most closely related to vom Lehn (2015), who also studies technological change and job polarization in a neoclassical framework. Nonetheless, our analysis differs along a number of dimensions, with notable distinctions in modeling framework, the measurement of automation technology, and the nature of the accounting experiment. We discuss these differences in detail in Section 6. Finally, we note that our work is also related to the recent contributions of Eden and Gaggl (2016) (who study the role of automation for trends in productivity and labor's share of national income) and Morin (2016) (who emphasizes the cyclical implications of computer adoption for routine employment).[3]

2. Empirical facts

We begin by documenting the decline in the share of the U.S population working in routine occupations. We analyze whether these changes are due to changes in the demographic composition of the economy, or to changes in the propensity to work in routine occupations conditional on demographic characteristics. We then identify specific demographic groups that account for the bulk of the changes in routine employment.

Our analysis uses data from the Monthly Current Population Survey (CPS), the main source of U.S. labor market statistics, made available through IPUMS (Flood et al., 2015). We focus on the civilian, non-institutionalized population aged 20–64 years old, excluding those employed in agriculture and resource occupations.[4] Following the literature (e.g. Acemoglu and Autor, 2011), we delineate occupations along two dimensions based on their task content: "cognitive" versus "manual," and

[2] See Autor and Dorn (2013) and Mazzolari and Ragusa (2013) who discuss the relation between the rise of non-routine manual employment and the decline in routine employment. With respect to the rise in non-employment in the U.S., Charles et al. (2013) discuss the role of the decline in manufacturing, Beaudry et al. (2016) highlight a reversal in the demand for cognitive skills, and Acemoglu et al. (2016) discuss the role of increased import competition.

[3] The empirical literature on job polarization is much larger and too vast to fully reference here. Much of the work exploits measures of "susceptibility" to automation based on the routine task intensity of employment (e.g. Autor and Dorn, 2013; Autor et al., 2015; Goos et al., 2014; Gregory et al., 2016); see (e.g. Michaels et al., 2014; Gaggl and Wright, 2016; Graetz and Michaels, 2016) for work using direct measures of ICT capital that largely focus on impacts at the industry or the firm level, rather than in aggregate.

[4] Given our interest in propensity changes for individuals with different levels of education, and the rise in educational attainment observed over time, one concern of including those in their early-20s might be misclassification due to college enrollment. In the Online Appendix (available as Supplementary material from the ScienceDirect website), we show that all results obtained here in Section 2 are essentially unchanged when we limit attention to 25–64 years old.

Table 1
Routine employment per capita.

	1979	1989	1999	2009	2014
Routine	0.405	0.406	0.376	0.317	0.312
Routine cognitive	0.173	0.196	0.182	0.169	0.161
Routine manual	0.232	0.210	0.194	0.148	0.151

Note: Share of the population employed in routine occupational groups based on individuals aged 20–64 from the monthly Current Population Survey, excluding those employed in agriculture and resource occupations.

"routine" versus "non-routine." The distinction between cognitive and manual occupations is based on the extent of mental versus physical activity. The distinction between routine and non-routine is based on the work of Autor et al. (2003). If the tasks involved can be summarized as a set of specific activities accomplished by following well-defined instructions, the occupation is considered routine. If instead the job requires flexibility, creativity, problem-solving, or human interaction, the occupation is non-routine. We group employed workers as either *non-routine cognitive*, *routine cognitive*, *routine manual* or *non-routine manual* based on an aggregation of 3-digit Census Occupation Codes. Details of the precise mapping are provided in Cortes et al. (2015). All statistics are weighted using person-level weights.

The decline in routine employment since the late 1980s has been well documented in the literature for many developed countries (e.g. Goos and Manning, 2007; Goos et al., 2009; Acemoglu and Autor, 2011; Jaimovich and Siu, 2012). Table 1 presents the population share of routine employment based on our CPS data. In 1979, routine occupations employed 40.5% of the working-age population in the U.S. This fraction remained stable over the following decade, and then began to decline steadily, until reaching a level of 31.2% in 2014. The breakdown between routine cognitive and routine manual employment reveals that the stability over the 1980s is due to offsetting changes in each of these groups. The share of the population employed in routine manual occupations declines over the entire 1979–2014 period. The population share of routine cognitive employment increases between 1979 and 1989, then declines until the end of the sample period. Given the different timing of the decline in routine manual and routine cognitive employment, we separately analyze the 1979–2014 and the 1989–2014 periods.

2.1. Decomposing labor market changes

The past four decades have also displayed marked changes in the educational and age composition of the population. Since demographic groups differ in their propensity to work in routine occupations, the decline of routine employment may be partially accounted for by these demographic changes. On the other hand, routine employment may be declining because of changes in the probability of working in such occupations for individuals with given demographic characteristics. These changes would be indicative of economic forces that change the labor market opportunities for specific groups of workers.

To investigate the relative importance of these two forces, we perform a set of decompositions where we divide the CPS sample into 24 demographic groups, based on the following three criteria: age, education, and gender. Specifically, we create three age groups (20–29, 30–49, 50–64, which we refer to as the young, prime-aged, and old respectively), four education groups (less than high school completion, high school diploma, some post-secondary, and college degree or higher), and two gender groups (females and males).

Denoting the fraction of the population in labor market state j at time t as $\overline{\pi}_t^j$, this can be written as:

$$\overline{\pi}_t^j = \sum_g w_{gt} \pi_{gt}^j, \tag{1}$$

where w_{gt} is the population share of demographic group g at time t, and π_{gt}^j is the fraction of individuals of demographic group g in state j at t. We consider five labor market states: employment in one of the four occupation groups described above, and non-employment (unemployment and labor force non-participation).

The change in the fraction of the population in state j can be written as:

$$\begin{aligned} \overline{\pi}_1^j - \overline{\pi}_0^j &= \sum_g w_{g1}\pi_{g1}^j - \sum_g w_{g0}\pi_{g0}^j \\ &= \sum_g \Delta w_{g1}\pi_{g0}^j + \sum_g w_{g0}\Delta\pi_{g1}^j + \sum_g \Delta w_{g1}\Delta\pi_{g1}^j. \end{aligned} \tag{2}$$

The first term, $\sum_g \Delta w_{g1}\pi_{g0}^j$, is a group size or *composition* effect, owing to the change in population share of demographic groups over time. The second component, $\sum_g w_{g0}\Delta\pi_{g1}^j$, is a *propensity* effect, due to changes in the fraction of individuals

Table 2
Decompositions based on age–education–gender groups.

	Pre	Post	Difference			
	(1)	(2)	Total (3)	Composition (4)	Propensity (5)	Interaction (6)
A. 1979–2014						
Number of Obs	*976,672*	*922,931*				
NRC (%)	21.5	28.2	+6.7	+9.7	−2.9	−0.0
RC (%)	17.3	16.1	−1.2	+0.6	−2.0	+0.3
RM (%)	23.2	15.1	−8.1	−5.2	−5.7	+2.7
NRM (%)	8.4	12.3	+3.9	−1.9	+6.6	−0.8
Not working (%)	29.6	28.3	−1.3	−3.1	+4.0	−2.2
B. 1989–2014						
Number of Obs	*977,282*	*922,931*				
NRC (%)	24.7	28.2	+3.5	+6.3	−2.7	−0.1
RC (%)	19.6	16.1	−3.5	+0.3	−3.9	+0.2
RM (%)	21.0	15.1	−5.9	−3.5	−4.0	+1.6
NRM (%)	9.6	12.3	+2.7	−1.7	+4.7	−0.3
Not working (%)	25.2	28.3	+3.1	−1.4	+5.9	−1.3

Notes: Composition of the population across different occupational groups and not working, based on individuals aged 20–64 from the monthly Current Population Survey, excluding those employed in agriculture and resource occupations. NRC stands for Non-Routine Cognitive, RC for Routine Cognitive, RM for Routine Manual, and NRM for Non-Routine Manual. Column (1) shows the composition for the initial period (1979 in Panel A; 1989 in Panel B); Column (2) shows the composition for the final period (2014 in both Panels). Column (3) shows the total change for the entire period, which is decomposed into the fraction attributable to changes in the composition of demographic groups in the population (Column (4)), changes in the propensity to enter the different categories conditional on demographic characteristics (Column (5)), and the interaction of the two (Column (6)). See text for full details.

within groups in state *j*. The third term, $sum_g \Delta w_{g1} \Delta \pi^j_{g1}$, is an *interaction* effect capturing the co-movement of changes in group sizes and changes in propensities.[5]

The results of this decomposition are presented in Table 2. We focus on two time intervals: the 35 year period from 1979 to 2014 exhibiting a monotonic decline in per capita employment in routine manual occupations in Panel A, and 1989–2014 exhibiting a similar decline in routine cognitive employment in Panel B. Columns (1) and (2) present the observed fraction of the population in each of the five labor market states. The total change in each of these population shares, displayed in Column (3), is decomposed into the composition, propensity, and interaction effect in Columns (4) through (6).

Panels A and B exhibit the well-documented increase in per capita employment in non-routine cognitive (NRC) occupations. This can be accounted for by composition change in the US population—the near doubling of the number of those with at least a college degree and, to a lesser extent, population aging (as both the highly educated and old have greater propensities for NRC work). In addition, we see increases in non-routine manual (NRM) employment in both periods, and a rise in non-employment during 1989–2014. Both of these are accounted for by propensity change, which we discuss further in Section 2.4.

Of principal interest are the declines in per capita routine manual (RM) employment in Panel A, and in per capita routine cognitive (RC) employment in Panel B. First, we note that the decline of RC is due entirely to declining propensities. In fact, the propensity change accounts for more than 100% of the total, as demographic change would have predicted an increase in the fraction of the population employed in RC occupations. With respect to the decline of RM employment, part of it is due to composition change, largely the shrinking population share with at most high school education. However, a greater proportion is due to the propensity effect—the fact that the likelihood of working in RM has fallen within demographic groups—either due to changes in behavior of otherwise identical individuals, or due to changing composition of unobservable characteristics for fixed demographic characteristics.[6] Indeed, given the strong increase in educational attainment, it is possible that the distribution of unobserved labor market productivity and/or leisure preferences of those with lower levels of education has shifted.

In the following subsections, we discuss how these composition and propensity changes for both routine occupational groups are concentrated within those of specific demographic characteristics.

[5] A common empirical approach to such accounting exercises is to perform a Oaxaca–Blinder (OB) decomposition (Blinder, 1973; Oaxaca, 1973). This would derive from a linear probability regression of inclusion in labor market states with age, education, and gender effects assumed to be additively separable. Our approach in Eq. (2) is equivalent to a OB specification where the regressors include a full set of interactions between demographic characteristics. This allows us to account for heterogeneity between groups in terms of propensity changes. Nonetheless, we display the results of the standard OB decomposition in the Online Appendix (available from the ScienceDirect website), and note that none of our findings are substantively altered.

[6] Note also that there is a partially offsetting interaction effect, implying that there is a positive correlation between the changes in group sizes and the changes in propensities.

Table 3.A
Fraction of change in Routine Manual employment accounted for by each demographic group, 1979–2014.

	Males			Females		
	20–29	*30–49*	*50–64*	*20–29*	*30–49*	*50–64*
Less than high school	**10.26**	**19.60**	**18.66**	3.60	8.41	5.60
High school diploma	**30.86**	**14.88**	−4.03	7.39	6.62	0.30
		All ages			*All ages*	
Some college		−13.55			−2.88	
At least college		−4.41			−1.33	

Notes: The table presents the fraction of the total change in the population share of Routine Manual (RM) employment that can be attributed to the changes experienced by each demographic group (by age, education and gender). The analysis is based on individuals aged 20–64 from the monthly Current Population Survey, excluding those employed in agriculture and resource occupations. The changes accounting for the majority of the total change are highlighted in bold.

Table 3.B
Key demographic groups: Routine Manual.

	Population share (%)			Fraction in RM (%)		
	1979	2014	Change	1979	2014	Change
Male high school dropouts						
Age 20–29	1.90	0.89	−1.01	61.58	37.87	−23.70
Age 30–49	4.12	2.06	−2.06	63.19	48.94	−14.25
Age 50–64	4.68	1.51	−3.17	43.09	32.92	−10.17
Male high school graduates						
Age 20–29	6.27	3.82	−2.45	61.36	34.99	−26.36
Age 30–49	7.51	6.60	−0.91	55.11	44.39	−10.72

Notes: The table presents the change in the population share and the propensity to be employed in RM occupations for these key groups.

2.2. Groups accounting for the decline in routine manual employment

To determine the importance of each demographic group in accounting for the decline in per capita RM employment, we compute the change induced by each group g, $w_{g1}\pi^j_{g1} - w_{g0}\pi^j_{g0}$ from Eq. (2), as a fraction of the total change.

The results are presented in Table 3.A. Five groups stand out as accounting for the bulk of the decline: male high school dropouts of all ages and male high school graduates under the age of 50. These groups combined can account for 94% of the fall in RM employment.

Table 3.B indicates that individuals of these demographic characteristics contribute to both the composition and propensity effects documented in Table 2. First, these groups are shrinking in terms of their share of the population (i.e., w_g is falling). While they represented nearly a quarter of the U.S. population in 1979, they represent less than 15% by 2014. Given that a large fraction of these low-educated men were employed in a routine manual occupation in 1979—as many as 63%, as indicated in the fourth column of the table—their fall in the population share has implied an important reduction in the overall share of RM employment, even holding propensities fixed.

More importantly, individuals within these key groups have experienced dramatic reductions in the propensity to work in RM (i.e., π_g is falling as well). For example, the fraction has fallen by about 25 percentage points for low-educated young men; while more than 60% of such individuals worked in a routine manual occupation in 1979, this is closer to one-third in 2014. As a result, the bulk of the propensity change documented in Table 2 is due to these five demographic groups.

Given that these key groups have experienced substantial movement out of RM employment, we ask where they have sorted into instead. We illustrate this in Table 4 by presenting the *change* in the share of each demographic group across labor market states. The results indicate that their dramatic decline in the probability of working in RM is offset primarily by increases in non-employment and, to a smaller extent, increases in non-routine manual employment. Clearly individuals of these demographic characteristics have not benefited from the increase in employment in high-paying, non-routine cognitive occupations observed in the aggregate.

2.3. Groups accounting for the decline in routine cognitive employment

We perform a similar analysis for the change in routine cognitive employment. Employment in these occupations peaks in 1989, so we focus on the 25 year period from then to the present. As documented in Table 2, more than 100% of the decline in per capita RC employment is accounted for by changes in propensity. Given this, we identify the key demographic groups in accounting for this propensity effect.

Table 4
Change in the fraction of workers in each group, 1979–2014 (p.p.).

	NRC	RC	RM	NRM	Not working
Male high school dropouts					
Age 20–29	−1.10	2.16	−23.70	7.47	15.17
Age 30–49	−4.95	0.62	−14.25	9.02	9.55
Age 50–64	−6.31	−0.12	−10.17	2.66	13.95
Male high school graduates					
Age 20–29	−3.81	5.22	−26.36	7.79	17.16
Age 30–49	−8.37	0.64	−10.72	5.32	13.13

Notes: The table details the changes in the fraction of workers in each occupational category and not working among the groups identified as accounting for the majority of the decline in RM employment. NRC stands for Non-Routine Cognitive, RC for Routine Cognitive, RM for Routine Manual, and NRM for Non-Routine Manual.

Table 5.A
Fraction of change in Routine Cognitive employment propensity accounted for by each demographic group, 1989–2014.

	Males			Females		
	20–29	*30–49*	*50–64*	*20–29*	*30–49*	*50–64*
High school diploma	−2.35	3.16	3.13	**14.80**	**24.13**	3.54
Some college	2.15	5.43	2.38	**12.27**	**10.62**	1.50
		All ages			*All ages*	
Less than high school		0.65			3.37	
At least college		8.75			6.46	

Notes: The table presents the fraction of the total change in the propensity to work in a Routine Cognitive (RC) occupation that can be attributed to the changes experienced by each demographic group (by age, education and gender). The analysis is based on individuals aged 20–64 from the monthly Current Population Survey, excluding those employed in agriculture and resource occupations. The changes accounting for the majority of the total change are highlighted in bold.

Table 5.B
Key demographic groups: Routine Cognitive.

	Population share (%)			Fraction in RC (%)		
	1989	2014	Change	1989	2014	Change
Female high school graduates						
Age 20–29	5.82	3.05	−2.77	32.61	22.73	−9.89
Age 30–49	10.58	5.57	−5.01	32.68	23.81	−8.87
Females with some college						
Age 20–29	3.88	4.70	0.82	36.77	24.46	−12.31
Age 30–49	5.48	6.32	0.84	33.04	25.50	−7.54

Notes: The table presents the change in the population share and the propensity to be employed in RC occupations for these key groups.

Table 5.A shows that the groups accounting for the bulk of the decline in RC propensity are young and prime-aged females with either high school diplomas or some post-secondary education. These four demographic groups alone account for 62% of the propensity effect.

The population shares and RC employment propensities for these groups are detailed in Table 5.B. All four groups experience obvious declines in their probability of working in RC, falling from approximately one-third in 1989 to one-quarter in 2014.

Given substantial movement out of RC employment, we ask where individuals with these characteristics have sorted into instead. Table 6 presents the change in the share of each demographic group across labor market states. As with the low-educated males identified in the decline of RM, these females with intermediate levels of education have not increased their propensity to work in high-paying, non-routine cognitive occupations. Instead, they have increased their propensities for non-employment and employment in non-routine manual occupations (with the former more prevalent among high school graduates, and the latter among those with some college). Relative to the males identified in the previous subsection, we generally observe smaller increases in non-employment rates among the female groups that account for the bulk in the decline in RC propensity.

Table 6
Change in the fraction of workers in each group, 1989–2014 (p.p.).

	NRC	RC	RM	NRM	Not working
Female high school graduates					
Age 20–29	−2.58	−9.89	−4.39	7.06	9.79
Age 30–49	−2.05	−8.87	−3.34	6.28	7.99
Females with some college					
Age 20–29	−4.42	−12.31	−1.16	9.94	7.96
Age 30–49	−3.78	−7.54	−0.24	7.44	4.11

Notes: The table details the changes in the fraction of workers in each occupational category and not working among the groups identified as accounting for the majority of the decline in the propensity to work in RC occupations. NRC stands for Non-Routine Cognitive, RC for Routine Cognitive, RM for Routine Manual, and NRM for Non-Routine Manual.

Table 7
Observed and counterfactual changes in population shares (p.p.).

	Observed (1)	Propensity (2)	Accounting CF (3)	Mitigating CF (4)
A. 1979–2014				
Routine	−9.30	−7.67	−6.20	−5.37
Non-routine manual	3.85	6.55	4.17	0.85
Non-employment	−1.27	4.03	3.14	−2.81
B. 1989–2014				
Routine	−9.37	−7.90	−5.68	−5.36
Non-routine manual	2.71	4.68	2.81	0.57
Non-employment	3.14	5.88	4.21	0.24

Column (1) shows the total observed change in the fraction of the population in different labor market categories, based on individuals aged 20–64 from the monthly Current Population Survey, excluding those employed in agriculture and resource occupations. Column (2) shows the counterfactual changes that are obtained when allowing for changes in the propensities to enter different labor market categories among all demographic groups, holding the composition of demographic groups in the population at benchmark levels. Column (3) shows the counterfactual changes that are obtained when holding the composition of all demographic groups in the population at benchmark levels, and holding the propensities at benchmark levels for all groups except those identified as being key for the decline in routine employment. Column (4) shows the counterfactual changes that are obtained when allowing the composition of demographic groups to change as in the data, while holding the propensities at benchmark levels only for the groups identified as being key for the decline in routine employment.

2.4. Aggregate importance of these demographic groups

As discussed above, the decline in per capita employment in routine occupations is due largely to declining probability to work in such occupations, as opposed to change in demographic composition, and the effect of declining propensity is concentrated in a subset of demographic groups. Here, we ask how much of the aggregate change in various labor market outcomes can be accounted for by the propensity change of these key demographic groups.

To determine this, we perform a number of simple counterfactual exercises in Table 7. The first column reproduces the change in the population share of routine employment, non-routine manual employment, and non-employment—the figures in Column (3) of Table 2. The second column reproduces the propensity effect from Column (5) of Table 2. Note that this represents a counterfactual holding the population shares of all demographic groups constant at their benchmark level (1979 in Panel A, 1989 in Panel B) and allowing *all* group-specific propensities to change as empirically observed.

The third column presents the result of a counterfactual in which only the propensities of the *key groups* are allowed to change; demographic composition and all other propensities are held constant at benchmark levels. This represents how much of the changes in Columns (1) and (2) are accounted for by the “behavioral changes” in our key groups. Of the approximate 9 percentage point fall in per capita routine employment displayed in either Panel A or B, about 65% is accounted for by the propensity change of our key groups; about three-quarters of the propensity effect in Column (2) is accounted for by the propensity effect of our key groups. This indicates the aggregate quantitative importance of the propensity change in the groups that we have identified.

Interestingly, even though the demographic groups were chosen based on their importance in accounting for the decline in routine employment, Table 7 shows that the behavioral change of these groups is also important in accounting for the aggregate changes in non-routine manual employment and non-employment. As evidenced in either panel, the propensity change of our key groups accounts for more than 100% of the observed increase in NRM employment, and about 60%

of the increase due to total propensity change. Similarly, these groups account for a large share of the increase in non-employment. The increase in the fraction not working is evident only in the 1989–2014 period. As Panel B indicates, the propensity change of our key groups accounts for more than 100% of the observed increase in non-employment, and about 70% of the propensity effect.

The fourth column presents a counterfactual in which demographic composition changes as observed in the US data, and all propensities change, except those of the key groups; these propensities are held constant at benchmark levels. This allows us to ask how much of the observed changes can be mitigated by omitting their behavioral change. As indicated in Panel A, if the propensity change of the key groups responsible for the decline of routine employment had not occurred, NRM employment would only have risen by 0.85 percentage points. This mitigates $3.00 \div 3.85 = 78\%$ of the observed increase. Similarly, in Panel B, omitting the key demographic groups mitigates $(3.14 - 0.24) \div 3.14 = 92\%$ of the observed increase in non-employment.

To summarize, the changes in employment and occupational choice of a small subset of demographic groups account for a large share of the decline in routine employment. These same groups are also key in understanding the rise of non-employment in the U.S. observed in the past 25 years and, to a slightly lesser extent, the rise of non-routine manual employment observed since 1979. This suggests that these long-run labor market changes are closely linked phenomena.

3. Model

Motivated by the findings of Section 2, we present a simple equilibrium model of the market for those low- and middle-skill workers identified as most responsible for the decline in routine employment over the past three decades (namely, young and prime-aged high school graduates, young and prime-aged females with some college education, and male high school dropouts of all ages). Our model is a generalized version of the model analyzed in Autor and Dorn (2013), extended along two key dimensions.

First, in addition to making an occupational choice between routine and non-routine manual jobs, individuals make a participation choice between working and non-employment. The empirical results presented in the previous section highlighted the importance of these two margins of labor adjustment. Second, we conduct our analysis making only minimal functional form and distributional assumptions on labor demand and labor supply. This generality allows us to characterize the theoretical and quantitative implications of progress in automation technology on labor market outcomes in a wide variety of parametric settings.

We use the model to study the role of advances in automation in rationalizing the changes in sorting of workers across employment in routine occupations, non-routine manual occupations, and non-employment. Given this goal, the analysis abstracts from other changes observed in the U.S. economy. For example, changes in the share of high-skilled workers and their occupational choice, changes in policy, and many other factors are likely to have contributed to the labor market outcomes discussed in Section 2. By concentrating solely on the impact of improvements in automation technology, we are able to present precise results from a general framework, and provide a template for further quantitative research in evaluating the role of automation.

3.1. Labor demand

Our theoretical results can be derived from a very agnostic specification of the demand for labor. In particular, we assume that GDP, Y_t, is produced with five factors of production via:

$$Y = G\left(K, L_C, L_M, \left[A + L_R^E\right]\right). \tag{3}$$

Here, K denotes capital (excluding the type of capital that relates to automation), L_C denotes the number of non-routine cognitive workers in the economy ("cognitive" hereafter), L_M denotes the number of non-routine manual workers ("manual"), L_R^E denotes the effective labor input of routine workers, and A denotes automation capital such as information and communication technology capital ("ICT" hereafter). As we discuss below, the amount of effective labor input differs from the measure of workers in the routine occupation. Effective routine labor and automation capital are assumed to be perfect substitutes in the production of "routine factor input" which we denote as $R = A + L_R^E$. This assumption allows the model to maximize the effect of automation on routine employment.

The representative firm hires factor inputs and sells output in competitive markets. Profit maximization results in demand for routine and manual labor that equates wages to marginal products:

$$W_R = G_R\left(K, L_C, L_M, \left[A + L_R^E\right]\right), \tag{4}$$

$$W_M = G_{L_M}\left(K, L_C, L_M, \left[A + L_R^E\right]\right). \tag{5}$$

Note that W_R denotes the wage per unit of effective labor.

3.2. Labor supply

Since the key demographic groups identified in Section 2 work almost exclusively in routine and manual occupations, we abstract from their ability to work in cognitive occupations. Hence, low-/middle-skilled individuals are assumed to make two discrete choices sequentially: first a decision whether to participate in employment or not and second, conditional on choosing to work, employment in the routine or manual occupation. We discuss these in reverse order.

Occupation decision: Individuals differ in their work ability, u, in the routine occupation where $u \sim \Gamma(u)$, where Γ denotes the cumulative distribution function (CDF). Given the wage per unit of effective labor, W_R, an individual (who has chosen to work) with ability u earns $u \times W_R$ if employed in the routine occupation. Alternatively, the worker earns W_M if employed in the manual occupation, independent of u (i.e., all low- and middle-skill workers have equal ability, normalized to 1, in manual work).

Denote by u^* the "cutoff ability level" such that individuals with $u < u^*$ optimally choose to work in the manual occupation, while those with $u \geq u^*$ choose the routine occupation. The cutoff is defined by the indifference condition:

$$u^* W_R = W_M, \tag{6}$$

for individuals who have chosen to participate in employment.

Participation decision: Individuals differ in their disutility of labor (or alternatively, their utility value of home production/leisure), b, where $b \sim \Omega(b)$, and Ω denotes the CDF. For simplicity, we assume that individuals choose whether to work prior to observing their routine work ability, u, knowing only that it is drawn from Γ.

As such, the expected return to working is given by:

$$b^* = W_M \Gamma(u^*) + W_R \int_{u^*}^{u^{max}} u \Gamma'(u) du \tag{7}$$

This anticipates the result that *ex post*, conditional on choosing to work, workers sort into the occupations according to the cutoff condition (6). Thus, *ex ante*, individuals with disutility $b < b^*$ choose to work, while those with $b \geq b^*$ optimally choose not to participate.[7]

3.3. Equilibrium

Labor market equilibrium implies that the demand for labor input in each occupation equals supply. Thus, for manual labor:

$$L_M = \Omega(b^*) \Gamma(u^*). \tag{8}$$

That is, given the participation rate, $\Omega(b^*)$, a fraction $\Gamma(u^*)$ of the workers work in the manual occupation. Similarly, the number of workers in the routine occupation is given by:

$$L_R = \Omega(b^*)[1 - \Gamma(u^*)]. \tag{9}$$

Finally, in terms of efficiency units, the effective routine labor input is given by:

$$L_R^E = \Omega(b^*) \int_{u^*}^{u^{max}} u \Gamma'(u) du. \tag{10}$$

3.4. The response to increased automation

Sections 4 and 6 study the response of the cutoff values u^* and b^*, which characterize sorting of workers across occupations and non-employment, to changes in capital-embodied automation technology. The six equations that will be used throughout the analysis are the two labor demand equations, (4) and (5), the two labor supply equations, (6) and (7), and two of the three market clearing conditions, (8) and (10).

We proceed by log-linearizing these equilibrium conditions. Denoting the percentage deviations of a variable from steady state by a circumflex, the demand for routine labor (4) becomes:

$$\widehat{W}_R = \eta_{G_R, L_M} \widehat{L}_M + \eta_{G_R, R} \left[\lambda \widehat{A} + (1 - \lambda) \widehat{L}_R^E \right], \tag{11}$$

where:

$$\lambda = \frac{A}{A + L_R^E} \in (0, 1). \tag{12}$$

Here, $\eta_{G_R,R}$ denotes the elasticity of the marginal product, G_R, with respect to the routine factor input, R, and η_{G_R,L_M} denotes the elasticity with respect to L_M. The log-linearization of the demand for manual labor (5) gives:

$$\widehat{W}_M = \eta_{G_{L_M}, L_M} \widehat{L}_M + \eta_{G_{L_M}, R} \left[\lambda \widehat{A} + (1 - \lambda) \widehat{L}_R^E \right], \tag{13}$$

[7] This sequential decision setup simplifies the model analysis. If individuals observed their disutility and routine work ability simultaneously, optimality would be characterized as a locus for the (b, u) cutoff.

where $\eta_{G_{L_M},L_M}$ is the elasticity of the marginal product, G_{L_M}, with respect to L_M and $\eta_{G_{L_M},R}$ is the elasticity with respect to the routine input, $R = A + L_R^E$.

The occupation choice condition (6) becomes:

$$\widehat{u}^* = \widehat{W}_M - \widehat{W}_R. \tag{14}$$

The log-linearization of the participation condition (7) implies:

$$b^*\widehat{b}^* = W_M\Gamma(u^*)\widehat{W}_M + W_M\Gamma'(u^*)u^*\widehat{u}^* + \left[W_R\int_{u^*}^{u^{max}} u\Gamma'(u)du\right]\widehat{W}_R - W_R\Gamma'(u^*)[u^*]^2\widehat{u}^*.$$

Using condition (6), this simplifies to become:

$$\widehat{b}^* = \psi\,\widehat{W}_M + (1-\psi)\widehat{W}_R, \tag{15}$$

where:

$$\psi = \frac{u^*\Gamma(u^*)}{u^*\Gamma(u^*) + \int_{u^*}^{u^{max}} u\Gamma'(u)du} \in (0,1). \tag{16}$$

Finally, the log-linearization of (8) and (10) imply:

$$\widehat{L}_M = \mu\,\widehat{b}^* + \nu\,\widehat{u}^*, \tag{17}$$

$$\widehat{L}_R^E = \mu\,\widehat{b}^* - \xi\,\widehat{u}^*, \tag{18}$$

where:

$$\mu = \frac{\Omega'(b^*)b^*}{\Omega(b^*)} \geq 0, \qquad \nu = \frac{\Gamma'(u^*)u^*}{\Gamma(u^*)} \geq 0, \tag{19}$$

and using Leibniz's rule:

$$\xi = \frac{\Gamma'(u^*)u^{*2}}{\int_{u^*}^{u^{max}} u\Gamma'(u)du} \geq 0. \tag{20}$$

Note that because $\Omega(b^*)$ is the employment participation rate, μ is the elasticity of the participation rate with respect to b^*. Similarly, since $\Gamma(u^*)$ is the fraction of workers who choose the manual occupation, ν is the elasticity of the "occupational choice" rate with respect to u^*. Finally, we note that ψ in Eq. (15) can be expressed as:

$$\psi = \frac{\xi}{\nu + \xi}. \tag{21}$$

Thus, the response of non-employment, routine employment, and manual employment, depends on parameters related to (i) the distribution of routine work ability, ν and ξ; (ii) the distribution of the disutility of labor, μ; (iii) the ratio of factors of production, λ; and (iv) own and cross elasticities of marginal products. The generality with which we have presented our framework allows the reader to simply "plug in" values of interest in order to evaluate the impact of changes in automation.

4. Theoretical analysis

In this section, we demonstrate the usefulness of the model of Section 3 in analyzing the role of progress in automation technology on labor market outcomes for the low-/middle-skilled workers of interest. We first determine the sign of the response of non-employment, routine employment, and manual employment to changes in automation technology, when imposing a minimal set of assumptions on model parameters. We then show how the presence or absence of a participation decision affects the response of sorting across routine and manual occupations (conditional on working). All proofs are presented in the Online Appendix (available as Supplementary material from the ScienceDirect website).

4.1. Signing the effects of automation

To proceed, we make the natural assumption that $\eta_{G_R,R} < 0$ and $\eta_{G_{L_M},L_M} < 0$; that is, production exhibits diminishing marginal product with respect to routine and manual factor inputs. We show by way of a simple example that the model is consistent with the empirical findings of Section 2, namely an increase in non-employment, $\widehat{b}^* < 0$, and an increase in manual versus routine employment (conditional on working), $\widehat{u}^* > 0$, in response to an increase in automation, i.e. $\widehat{A} > 0$.

Proposition 1. *Let the cross elasticities in production be zero, i.e.* $\eta_{G_R,L_M} = \eta_{G_{L_M},R} = 0$. *Then, for all values of* λ, μ, ν, *and* ξ, *an increase in automation technology increases non-employment, and reallocates employment from the routine to the manual occupation.*

The economics of this case are as follows. In response to an increase in automation, the supply of routine factor input increases. Given diminishing returns, this leads to a fall in the routine occupation wage. Since cross elasticities are zero, the wage in the manual occupation is not affected directly by the change in automation. As such, conditional on participation, workers move from the routine to the manual occupation. Given diminishing (or even constant) marginal product of manual labor, the wage in the manual occupation is either falling (or constant). Since the return to employment is a weighted average of the routine and manual wages, the *ex ante* wage falls. Hence, participation falls.

4.2. The effects of a participation margin

As discussed above, an increase in automation causes workers to leave the routine occupation and sort into the manual occupation. Here we explore how the inclusion or exclusion of an employment participation choice affects the degree of occupational reallocation. Since we are especially interested in the case when the degree of occupational reallocation is maximized, we assume non-diminishing marginal product of manual labor, $\eta_{G_{L_M},L_M} = 0$. From Eq. (13), this eliminates the fall in the manual wage as workers move into the manual occupation. We obtain the following result.

Proposition 2. *Let the following elasticities in production be zero:* $\eta_{G_R,L_M} = \eta_{G_{L_M},R} = \eta_{G_{L_M},L_M} = 0$. *Then, the presence of a participation margin mitigates the degree of occupational reallocation from the routine to the manual occupation.*

As we show in the Online Appendix, the response of the occupation cutoff to an automation shock is given by:

$$\widehat{u}^* = \left[\frac{\lambda}{(1-\lambda)\xi - \frac{1}{\eta_{G_R,R}} + (1-\psi)(1-\lambda)\mu}\right]\widehat{A}. \tag{22}$$

Note the final term in the denominator, $(1-\psi)(1-\lambda)\mu \geq 0$, and recall from Eqs. (17) and (18) that μ is the elasticity of the participation rate with respect to b^*. Hence, all else equal, the response of occupational reallocation, $\widehat{u}^*$, is maximized when participation does not adjust, $\mu = 0$. In other words, the higher is the elasticity parameter on participation, the smaller is the movement of workers from routine to manual (conditional on working).

The intuition is as follows. An increase in automation technology drives down the marginal product of routine labor. With a constant wage in the manual occupation, there must be movement of labor out of the routine occupation, until the cutoff Eq. (6) is satisfied for the marginal worker. When there is no employment participation choice, the number of workers is fixed; all adjustment comes from occupational reallocation out of routine jobs. With endogenous participation, some of the adjustment comes from fewer workers selecting into employment. All else equal, a reduction in the number of workers raises the marginal product of routine labor; this allows equilibrium to be attained with less occupational reallocation than otherwise.

In summary, this proposition highlights an important tradeoff in neoclassical analyses of automation's impact on labor market outcomes. Maximizing the impact of advances in automation on occupational reallocation requires abstracting from participation choice. However, endogenizing the decision to work mitigates the impact on occupational sorting.

5. Quantitative specification

In the next two sections we study the quantitative effect of automation. We discuss the quantitative specification of the model here. We pick 1989 as the "steady state" around which the model economy is linearized, since this year corresponds to the maximum in per capita routine employment as displayed in Table 1. Numerical results are presented in Section 6.

5.1. Shares

Labor values: While the demographic groups identified in Section 2 account for the bulk of routine employment, other groups account for important shares of manual employment and non-employment. The Online Appendix details the shares (of population, employment, income) delineated across the "key" and "other" groups used in the quantitative analysis, and we highlight selected values here that are relevant for the analysis.

Within the key group, the employment rate (or employment "propensity" using the terminology of Section 2) was 72.7% in 1989, and conditional on employment, the fraction working in routine occupations was 81.6% (to align with the model, we exclude the 13.7% working in cognitive occupations). Given a specification of the distribution of routine work ability (discussed below), we calculate the average efficiency in routine work for our key group. We use this average efficiency as the efficiency of routine workers from the other demographic group. This allows us to have a simple aggregation of the effective routine labor input.

As a point of reference in evaluating the results of Section 6, the employment rate of the key group fell to 64.9% by 2014, and the fraction of workers employed in routine occupations fell to 69.1%.

Routine factor inputs: Calibrating λ in Eq. (11) requires specifying the ratio of service flows from automation capital, A, to effective routine labor, L_R^E. Since we have specified these inputs to be perfect substitutes in production, we can measure this ratio empirically as the ratio of their factor shares of national income. Using CPS data and data from Eden and Gaggl (2016) (see below) we obtain a 1989 value of $\lambda = 0.0845$.

5.2. Elasticities and distributions

Participation elasticity: Recall that μ is the elasticity of the participation rate (in employment versus non-employment) with respect to b^*. To quantify this participation elasticity, we decompose the elasticity of the participation rate with respect to b^* into: (i) the elasticity of the participation rate with respect to the wage, divided by (ii) the elasticity of b^* with respect to the wage.

Part (i) can be identified empirically. The literature studying the earned income tax credit (EITC) provides various estimates of the wage elasticity of the participation rate. The handbook chapter by Hotz and Scholz (2003) suggests a value between 0.97 and 1.69. As such we take 1.3 as a benchmark estimate. Part (ii) can be pinned down theoretically. The cutoff condition, Eq. (15), implies that the elasticity of b^* with respect to an equal percentage change in the routine and manual wage equals 1. Hence, we specify $\mu = 1.3$ as a useful benchmark in our numerical analysis. However, given that the mapping between the EITC literature and our model is not perfect, we consider also a higher value of $\mu = 2$.[8]

Production function elasticities: We restrict attention to the case when own elasticities are negative, $\eta_{G_R,R} < 0$ and $\eta_{G_{L_M},L_M} < 0$. With respect to cross elasticities in production, we consider several cases. First, we study cross elasticities that are zero, i.e. $\eta_{G_R,L_M} = \eta_{G_{L_M},R} = 0$.

For cases where cross elasticities differ from zero, we note that our neoclassical production framework implies the following relation across the different production elasticities

$$\frac{(W_M L_M)/Y}{(W_R R)/Y} \times \left(\eta_{G_{L_M},R}\right)^2 = \eta_{G_R,R} \times \eta_{G_{L_M},L_M}. \tag{23}$$

There are two steps to this derivation. First, note that the cross elasticities are given by

$$\eta_{G_R,L_M} \equiv \left(\frac{\partial G_R}{\partial L_M}\right)\left(\frac{L_M}{G_R}\right), \tag{24}$$

$$\eta_{G_{L_M},R} \equiv \left(\frac{\partial G_{L_M}}{\partial R}\right)\left(\frac{R}{G_{L_M}}\right), \tag{25}$$

which using the fact that the firm's FOCs (4) and (5) require wages to equal marginal products, can be written as

$$\eta_{G_R,L_M} \equiv \left(\frac{\partial G_R}{\partial L_M}\right)\left(\frac{L_M}{W_R}\right), \tag{26}$$

$$\eta_{G_{L_M},R} \equiv \left(\frac{\partial G_{L_M}}{\partial R}\right)\left(\frac{R}{W_M}\right), \tag{27}$$

Using Young's theorem, $\partial G_R/\partial L_M = \partial G_{L_M}/\partial R$, it follows that the ratio of the cross elasticities is given by:

$$\frac{\eta_{G_R,L_M}}{\eta_{G_{L_M},R}} = \frac{W_M L_M}{W_R R} = \frac{(W_M L_M)/Y}{(W_R R)/Y}. \tag{28}$$

Hence, the ratio of elasticities must equal the ratio of manual labor's share of income to the share of income paid to all routine factors of production.[9] In the data, this ratio of income shares is equal to 0.1355, disciplining the relative magnitude of cross elasticities.

The second step is that by definition the product of the cross elasticities must equal the product of the own elasticities, i.e.

$$\eta_{G_R,L_M} \times \eta_{G_{L_M},R} \equiv \eta_{G_R,R} \times \eta_{G_{L_M},L_M}. \tag{29}$$

Combining then Eqs. (28) and (29) leads to (23). Eq. (23) provides an empirical restriction on the (absolute value of the) cross elasticities as a function of the own elasticities, while permitting routine and manual inputs to be either gross complements or substitutes. In considering different values of the own elasticities, we refer to the empirical work of Lichter et al. (2015) who conduct a meta analysis of 151 different studies containing 1334 estimate of the own-wage elasticity of labor demand. We consider the range of estimates that they provide for the U.S. which lies between 0 and −3.[10] Then, for every combination of the own elasticities we use (23) to recover the value for $\eta_{G_{L_M},R}$. Using this value of $\eta_{G_{L_M},R}$ in (28) we recover the value of η_{G_R,L_M}.

Routine ability distribution: Finally, to specify ν and ξ, we need to make choices on the distribution of routine work ability, Γ. We consider three possibilities to explore the robustness of our findings. The first is a degenerate distribution of

---

8 As will become clear, values lower than $\mu = 1.3$ make the model unable to match the facts regarding participation.

9 This uses the fact that the wage per unit of effective routine labor must equal the rental rate per unit of automation capital service flow in equilibrium.

10 Lichter et al. (2015) have assembled 287 estimates for the U.S. Of these, 11 are positive which we discard from the analysis. Of the remaining 276 estimates, about 95% lie between 0 and −3 which is the range we consider. We are grateful to the authors of Lichter et al. (2015) for kindly sharing their data with us; all errors in their use are our own.

routine ability, equal to the ability in the manual occupation. In Section 6, we discuss how assuming identical work ability in both occupations generates sharp results that are *independent* of the production elasticities.

In the remaining two cases, we assume Γ to be either uniform or Pareto. In the case of the uniform skill distribution, we specify the support to be $[0, u^{max}]$ so that:

$$\Gamma(u) = \frac{u}{u^{max}}. \tag{30}$$

This implies:

$$\xi = \frac{u^{max} u^{*2}}{u^{max} - u^{*}}, \qquad \nu = 1. \tag{31}$$

In the case of the Pareto distribution:

$$\Gamma(u^{*}) = 1 - \left(\frac{u_{min}}{u^{*}}\right)^{\kappa_u}, \tag{32}$$

which implies:

$$\xi = \kappa_u - 1, \qquad \nu = \kappa_u \left(\frac{u_{min}}{u^{*}}\right)^{\kappa_u} \left[1 - \left(\frac{u_{min}}{u^{*}}\right)^{\kappa_u}\right]^{-1}. \tag{33}$$

In both cases, we have two values to pin down in order to specify ν and ξ. These are u^{max} and u^{*} for the uniform, κ_u and (u_{min}/u^{*}) for the Pareto. We use the same two moments in the data to identify these. The first is the fraction of workers (conditional on working) employed in the manual occupation in 1989, $\Gamma(u^{*}) = 0.184$. The second data moment is the ratio of national income shares paid to routine and manual workers in 1989; the ratio of income shares is given by $\left(\frac{u^{max} - u^{*}}{u^{max} u^{*2}}\right)$ for the uniform distribution and $\left(\frac{\kappa_u}{\kappa_u - 1}\right)\left(\frac{u_{min}}{u^{*}}\right)^{\kappa_u}\left[1 - \left(\frac{u_{min}}{u^{*}}\right)^{\kappa_u}\right]^{-1}$ for the Pareto distribution. We measure this ratio to equal 6.757. Solving these two moment conditions for two unknowns results in $\xi = 0.148$ and $\nu = 1$ for the uniform distribution, and $\xi = 1.680$ and $\nu = 11.866$ for the Pareto distribution.[11]

5.3. Automation shock

To measure the increase in automation since 1989 (or the magnitude of the "automation shock," given our focus on linearized dynamics), we relate capital-embodied automation technology, A, to measured ICT capital using the data of Eden and Gaggl (2016). Crucially, ICT capital is expressed in real, *effective* units in their data, using detailed asset-level price deflators for both investment flows and existing stocks, obtained from the BEA Fixed Asset Accounts. This aligns Eden and Gaggl (2016)'s empirical measure with the construct of the model.

Simply using the percentage growth rate of ICT capital between 1989 and 2014 neglects the fact that along a balanced growth path (BGP) with constant labor allocations, all forms of capital are expected to grow, at potentially different rates. As is well known, the growth rate of ICT capital has exceeded that of other forms of capital at least since the 1970s, prior to the decline of per capita routine employment, which as we discuss below is consistent with a BGP. As such, we measure the shock as the *deviation* of ICT capital from a balanced growth trend with "automation-specific" technical change since 1989.

To do so, we adapt the now-standard methodology of Greenwood et al. (1997) to the model of Section 3. The details are provided in the Online Appendix. Briefly, we combine the price and quantity data of Eden and Gaggl (2016) with NIPA data for Real GDP to construct a counterfactual series for the stock of automation capital that would have obtained had the economy been along a BGP from 1989 to 2014. We then compare the actual series to the counterfactual, and find the actual series to be approximately 100 log points higher in 2014. Thus, in our quantitative analysis we use $\widehat{A} = 1$ as our automation shock and consider robustness to alternative values.

6. Numerical results

Here we evaluate the model's quantitative predictions. As the previous section indicates, our analysis is broad and agnostic, within the neoclassical tradition of Greenwood et al. (1997) and Krusell et al. (2000)—frameworks that have become standard building blocks to analyze the effects of capital-embodied technical change and its labor market implications. Our approach allows for an agnostic specification of labor demand with few restrictions, and a standard approach to labor supply with respect to employment and occupational choice. But obviously, the results rely on specific qualitative and quantitative choices. For example, automation is proxied as efficiency units of ICT capital, and advances in automation measured as deviations of its stock from a (specific) balanced growth trend. Different assumptions would lead to different results, warranting future investigation that we believe is aided by the simple, linearized analytical framework developed here.

Given an increase in capital-embodied automation technology as measured in Section 5, we solve for the response of sorting between routine/manual occupations and non-employment choice for the key demographic group. Recall that in the

[11] For the Uniform distribution ξ also equals the ratio of the manual to routine workers, hence the value of $0.148 = 6.7568^{-1}$.

data, the key group's employment rate fell from 72.7% to 64.9%, and conditional on employment, the fraction working in manual occupations rose from 18.4% to 30.9%, between 1989 and 2014. We begin the analysis with the case of a degenerate distribution of routine work ability, equal to the ability in the manual occupation. This serves as a useful benchmark.

6.1. Homogeneous routine ability

In this case all workers are equally productive in both occupations (with productivity normalized to unity). Thus, in equilibrium $W_M = W_R$; furthermore, the participation equation implies that $b = W_M = W_R$. In the Online Appendix, we show that *irrespective* of the production elasticities, following an automation shock, the response of the fraction of workers (conditional on working) who sort into the manual occupation is given by:

$$(1 - \Gamma(u^*))\left(\frac{\lambda}{1-\lambda}\right)\widehat{A}, \tag{34}$$

where, with a slight abuse of notation, $(1 - \Gamma(u^*))$ denotes the steady state fraction of workers who sort into the routine occupation in 1989. Furthermore, the response of the employment rate is given by:

$$-(1 - \Gamma(u^*))\left(\frac{\lambda}{1-\lambda}\right)\widehat{A}. \tag{35}$$

This allows for an easy quantification of the effects of automation. From Section 5, $\lambda/(1-\lambda) = 0.0923$. The fraction of workers in the key group employed in manual occupations in 1989 is 0.184. Given our estimate of $\widehat{A} = 1$, following the automation shock the fraction of workers employed in manual occupations equals $\exp[0.0923 \times (1 - 0.184) \times 1 + \log(0.184)] = 0.198$. Similarly, the employment rate falls to $\exp[-0.0923 \times (1 - 0.184) \times 1 + \log(0.727)] = 0.674$. The corresponding empirical values in 2014 are 30.9% and 64.9%, respectively. Hence, the model underpredicts both the reallocation to the manual occupation and the reallocation to non-employment.

Finally, the simplicity of this analysis also allows us to ask what value of $\widehat{A}$ is required to "reverse engineer" the changes observed in the data for the key group. To account for all of the occupational reallocation via the automation shock would require $\widehat{A} = 6.04$. This value would also allow the model to explain more than 100% of the change in employment rate. Hence, the change in automation technology, in log deviation terms, would need to be *six* times greater than the one indicated in Section 5.

6.2. Heterogeneous routine ability

In this subsection, we analyze the case with heterogeneous routine work ability when cross elasticities in production are zero, i.e. $\eta_{G_R,L_M} = \eta_{G_{L_M},R} = 0$. This implies that changes in automation do not have a direct effect on the marginal product of manual labor. This case is informative given the analytical results of Propositions 1 and 2 presented in Section 4.

In the rest of this Section, we discuss the specification with Pareto distributed ability; the results with the uniform distribution are presented in the Online Appendix. We solve the model for values of the own elasticities, $\eta_{G_R,R}$ and $\eta_{G_{L_M},L_M}$, that lie in the interval $[-3, 0)$, with the elasticity of participation set to the benchmark value of $\mu = 1.3$. Recall that:

$$\Omega(b^*) = \text{Participation rate}, \tag{36}$$

$$\Gamma(u^*) = \text{Fraction in manual occupation}, \tag{37}$$

implying that in the linearized equilibrium:

$$\frac{\Omega'(b^*)b^*}{\Omega(b^*)}\widehat{b} = \mu \times \widehat{b} = \text{Partici}\widehat{\text{pation}}\text{ rate}, \tag{38}$$

$$\frac{\Gamma'(u^*)u^*}{\Gamma(u^*)}\widehat{u} = \nu \times \widehat{u} = \text{Fraction in }\widehat{\text{manual}}\text{ occupation}. \tag{39}$$

For each combination of parameter values, the solution of the linearized model yields values for $\widehat{u}$ and $\widehat{b}$, from which we recover the employment rate and occupational sorting in response to the automation shock.

Each dot in Fig. 1 depicts the employment rate on the vertical axis and, conditional on working, the fraction of workers in manual occupations on the horizontal axis, for specific pairs of $(\eta_{G_R,R}, \eta_{G_{L_M},L_M})$. These elasticities are reported, in that order, to the right of each dot (for visual clarity we do so only for selected elasticity pairs).

Fig. 1 illustrates that the effects of an increase in automation are of the correct sign relative to Proposition 1: with the original 1989 allocation located at the upper-left corner, employment falls and workers reallocate toward the manual occupation. But matching the magnitude of change in either variable is a challenge for the model. In terms of occupational reallocation, the model can account for at most about 50%, represented by the point furthest to the right in the diagram. In this parametrization, the model accounts for less than one-third of the fall in the employment rate. This occurs when $\eta_{G_{L_M},L_M} = -0.001$ and $\eta_{G_R,R} = -3$, when the demand curve for manual labor is flat, and much steeper for the routine factor input. Thus, the elasticities must be at opposing extreme values.

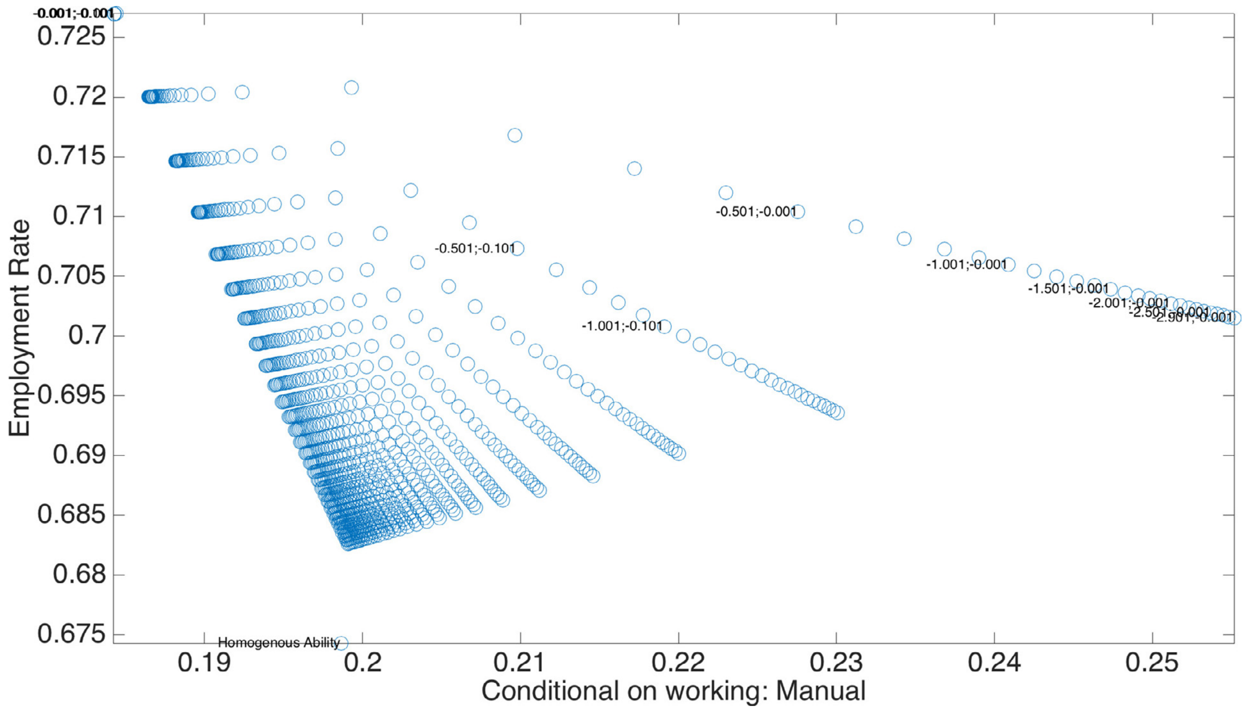

Fig. 1. Employment rate and occupation choice. *Notes:* Each dot corresponds to different values of $(\eta_{G_R,R}, \eta_{G_{L_M},L_M})$ as indicated to the right of selected dots. All simulations use a Pareto distribution for ability and $\mu = 1.3$. In 1989, the employment rate was 72.7%, and conditional on employment, the fraction working in manual occupations was 18.4% ("north–west" corner). In 2014 the employment rate fell to 64.9%, and conditional on employment, the fraction in manual occupations increased to 30.9%, off the scale of the figure in the "south–east" direction.

When the elasticities are flipped, with $\eta_{G_{L_M},L_M} = -3$ and $\eta_{G_R,R} = -0.001$, the model generates a larger fall in employment, but very little reallocation of labor across occupations. This highlights the importance of the relative magnitudes of $\eta_{G_R,R}$ and $\eta_{G_{L_M},L_M}$. Consider the "flattest" locus of points with $\eta_{G_{L_M},L_M} = -0.001$ (the second number to the right of each dot). As $\eta_{G_R,R}$ becomes more negative, there is greater response of occupational reallocation in response to an increase in automation; however, there is comparatively little change in the employment rate. Larger responses in employment require values of $\eta_{G_{L_M},L_M}$ that are larger in absolute value (more negative). But along any locus where $\eta_{G_R,R}$ is constant (the first number to the right of a dot), increasingly negative values of $\eta_{G_{L_M},L_M}$ move in the "wrong" south-westerly direction.

Hence, Fig. 1 illustrates a quantitative tradeoff between responsiveness on the participation and occupational sorting margins. Greater occupational reallocation requires values of $\eta_{G_{L_M},L_M}$ closer to zero, so that the manual wage does not fall "too fast" with increased employment. But this relatively flat labor demand curve implies little change in the employment rate. Generating greater responses of employment to automation requires steep labor demand curves, resulting in little response in occupational sorting.

Fig. 1 is also useful to explore the effect of heterogeneity on our results. Near the bottom-left corner of the figure, we plot the response generated from the version with homogenous work ability across occupations discussed above. For every point with $\eta_{G_R,R} < \eta_{G_{L_M},L_M}$ the amount of occupation reallocation is higher in the presence of heterogeneity. This can be shown formally in the limiting case for $\eta_{G_R,R} < \eta_{G_{L_M},L_M} = 0$.[12] In particular, as workers move from routine to manual occupations, the routine wage rises. The greater the heterogeneity in routine ability, the smaller is the impact of the marginal worker on effective routine labor input which determines the wage (marginal product). Thus, with greater heterogeneity, more reallocation of workers out of the routine occupation is required to satisfy the occupational sorting condition, Eq. (14).

This analysis also has implications for the behavior of relative wages. Eq. (14) (in linearized form, which we reproduce here):

$$\widehat{u}^* = \widehat{W}_M - \widehat{W}_R, \tag{40}$$

[12] Specifically, with homogenous ability, the change in the number of workers working in the manual occupation is given by $\left(\frac{1-\Gamma(u^*)}{\Gamma(u^*)}\right)\left(\frac{\lambda}{1-\lambda}\right)\widehat{A}$. In the case with Pareto distributed ability, the change is given by $\left(\frac{\kappa_u}{\kappa_u-1}\right)\left(\frac{1-\Gamma(u^*)}{\Gamma(u^*)}\right)\left(\frac{\lambda}{1-\lambda}\right)\widehat{A}$. Greater heterogeneity (i.e. a thicker right tail, $\kappa_u \to 1$) increases the reallocation from routine to manual; as $\kappa_u \to \infty$ the Pareto converges to a degenerate distribution, and the response of occupation sorting converges to that with homogenous ability.

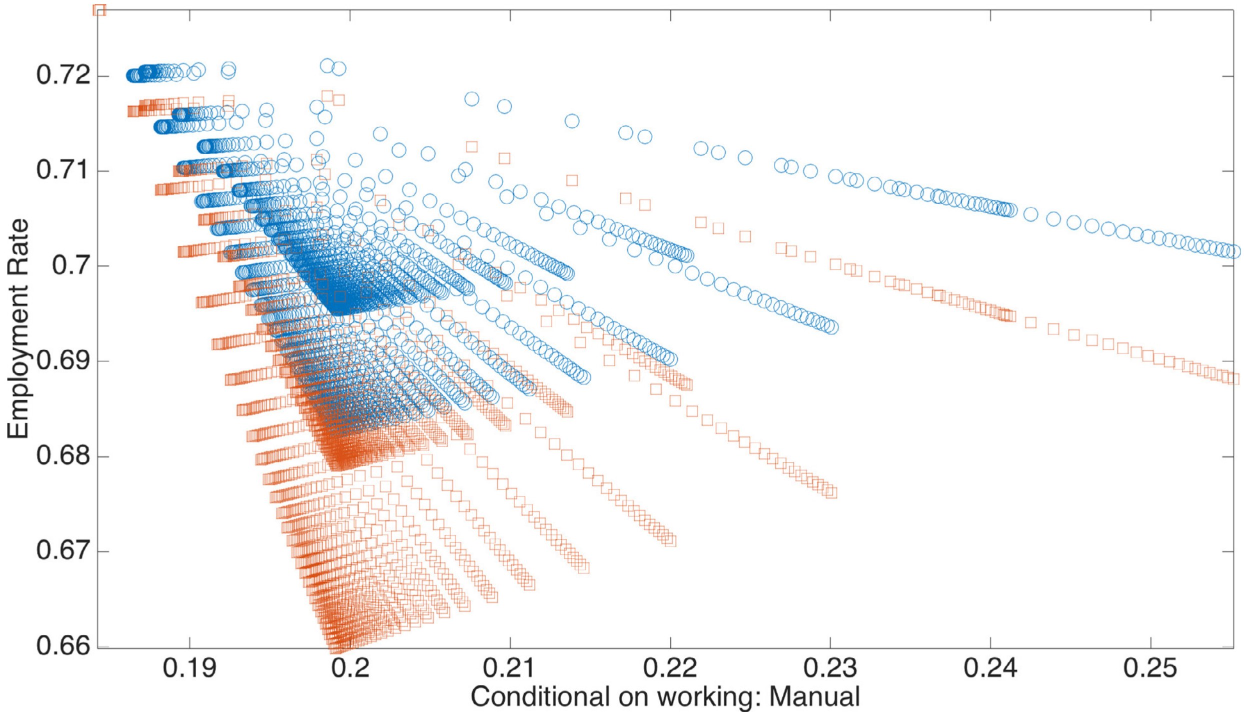

Fig. 2. Employment rate and occupation choice: the effect of μ. *Notes:* Each dot corresponds to different values of $(\eta_{G_R,R}, \eta_{G_{L_M},L_M})$ and μ. The blue circle dots use $\mu = 1.3$, the red square dots use $\mu = 2$. All simulations use a Pareto distribution for ability. In 1989, the employment rate was 72.7%, and conditional on employment, the fraction working in manual occupations was 18.4% ("north–west" corner). In 2014 the employment rate fell to 64.9%, and conditional on employment, the fraction in manual occupations increased to 30.9%, off the scale of the figure in the "south–east" direction. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

states that the change in the marginal ability, u^*, is equal to the change in relative wages (per effective unit of labor). As we show in the Online Appendix, the ratio of manual to routine wage per effective unit increased by about 10%, 1989–2014.

This implies an increase in u^* and, thus, an increase in the fraction of workers who select into the manual occupation. All cases in Fig. 1 are qualitatively consistent with this. Quantitatively, note that the maximal change in u^* is $\widehat{u}^* = 0.0360$ or 3.6%, corresponding to the maximal value obtained on the horizontal axis. This is only about one-third of the empirically observed change in relative wages, another manifestation of the challenge that the benchmark model faces in rationalizing the data.

Fig. 2 investigates the effect of changes in the participation elasticity, with $\mu = 1.3$ depicted as circle dots, and $\mu = 2$ as square dots. For any pair of the production elasticities, the larger is μ the greater is the response of the employment rate. However, for the square dot furthest to the right in the diagram, the model accounts for only about 50% of the change in both occupational sorting and the fall in employment.[13]

The figures in the Online Appendix replace the analysis of Fig. 1 for the case when routine work ability is uniformly distributed. As in the Pareto case, the overall impact of an increase in automation on labor market outcomes is small. Moreover, the effect on occupational sorting is smaller than in the Pareto case displayed above.[14]

To summarize, when changes in automation have no direct effect on the marginal product of labor in the manual occupation (i.e., when cross elasticities are zero), this neoclassical framework has only modest ability to generate the response of employment and occupational sorting observed in the data. This is true despite considering a wide range of parameter values.

6.3. Non-zero cross elasticities

We now discuss results when cross production elasticities differ from zero. The empirical literature is silent on whether manual and routine inputs are gross complements or substitutes. With respect to analytical results, Autor and Dorn (2013) show that for their functional form assumptions, simultaneous employment and wage polarization can only

[13] In experiments not presented here, results were largely unchanged for greater values of μ that far exceed those in the range reported in Hotz and Scholz (2003).

[14] This can be formally shown for the limiting case of $\eta_{G_{L_M},L_M} = 0$ and $\eta_{G_R,R} \to -\infty$.

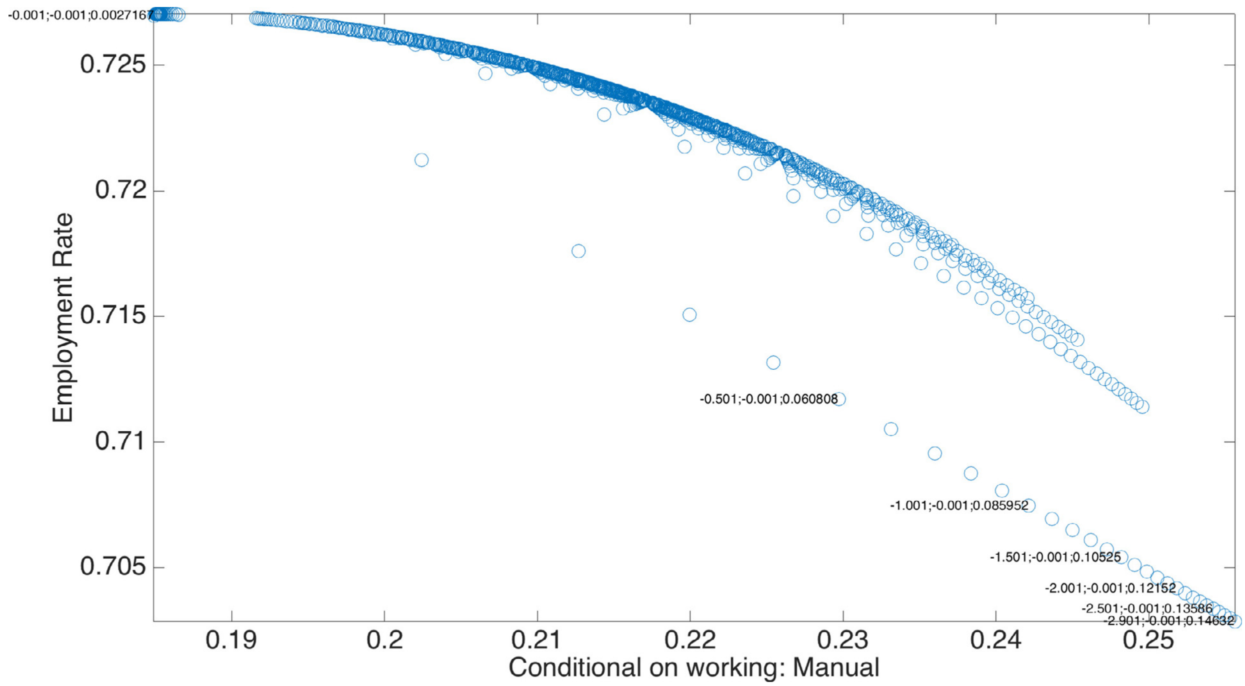

Fig. 3. Employment rate and occupation choice: complements. *Notes:* Each dot corresponds to different values of $(\eta_{G_R,R}, \eta_{G_{L_M},L_M}, \eta_{G_{L_M},R})$ as indicated to the right of selected dots. All simulations use a Pareto distribution for ability and $\mu = 1.3$. In 1989, the employment rate was 72.7%, and conditional on employment, the fraction working in manual occupations was 18.4% ("north–west" corner). In 2014 the employment rate fell to 64.9%, and conditional on employment, the fraction in manual occupations increased to 30.9%, off the scale of the figure in the "south–east" direction.

be rationalized when manual and routine labor are gross complements. For the sake of completeness and generality, we consider both positive and negative values of cross elasticities.

Fig. 3 depicts the case when the two occupational inputs are complements, Fig. 4 for the case of substitutes. Each dot depicts the employment rate on the vertical axis, and the fraction employed in the manual occupation (conditional on working) on the horizontal axis, for specific triplets of $\eta_{G_R,R}$, $\eta_{G_{L_M},L_M}$, $\eta_{G_{L_M},R}$ (recall that η_{G_R,L_M} is determined by equation (23)). For visual clarity, we report the elasticity values (in the listed order, to the right of each dot) only for the case when $\eta_{G_{L_M},L_M} = -0.001$, but the figures report all the different elasticity combinations. Again, the Online Appendix displays the results for the case when routine work ability is uniformly distributed.

In this case, the automation shock has a direct effect of increasing the marginal product of labor in the manual occupation. This amplifies reallocation towards the manual occupation. However, it severely reduces the model's ability to generate reductions in employment, since the negative effect of the shock on wages is dampened. Accordingly, Fig. 3 displays a larger set of parameters with more reallocation towards the manual occupation (conditional on working) and less change in the employment rate relative to Fig. 1. Overall, the quantitative implications of the complements case are similar to the case when the cross elasticities are zero and $\eta_{G_{L_M},L_M} = -0.001$.

When the two inputs are substitutes, the automation shock has a direct effect of decreasing the marginal product of labor in the manual occupation. Not surprisingly, Fig. 4 indicates that this dampens the incentive to reallocate towards the manual occupation. In fact, there are now parameterizations where an increase in automation causes reallocation *towards* the routine occupation, conditional on working. Overall, considering cross elasticities in production that differ from zero does not substantively alter the conclusion from Sections 6.1 and 6.2. The neoclassical framework is only modestly successful at generating the observed changes in the low- and middle-skill labor market in response to increased automation.

Our findings contrast with those in vom Lehn (2015). We note three important distinctions. The first is the nature of the analysis; while we focus *solely* on quantifying the role of advances in automation technology, vom Lehn (2015), considers *both* the role of technical change and changes in the distribution of labor market skill. A second difference is in terms of the measurement of capital-embodied technology; we equate automation with ICT capital, whereas vom Lehn (2015) considers total non-residential equipment. Finally, we consider labor market responses to deviations in the *empirically observed* ICT capital stock from a BGP. In contrast, vom Lehn (2015) studies a representative agent problem for consumption, investment, and hours worked given a time path of *equipment investment specific technical change*. As such, the resulting capital stock in vom Lehn (2015) is the one implied by his model dynamics, as opposed to actual changes in observable capital stocks which we use. Hence, our analysis is more prescriptive in nature, providing an agnostic framework for evaluating results for a wide range of parameterizations.

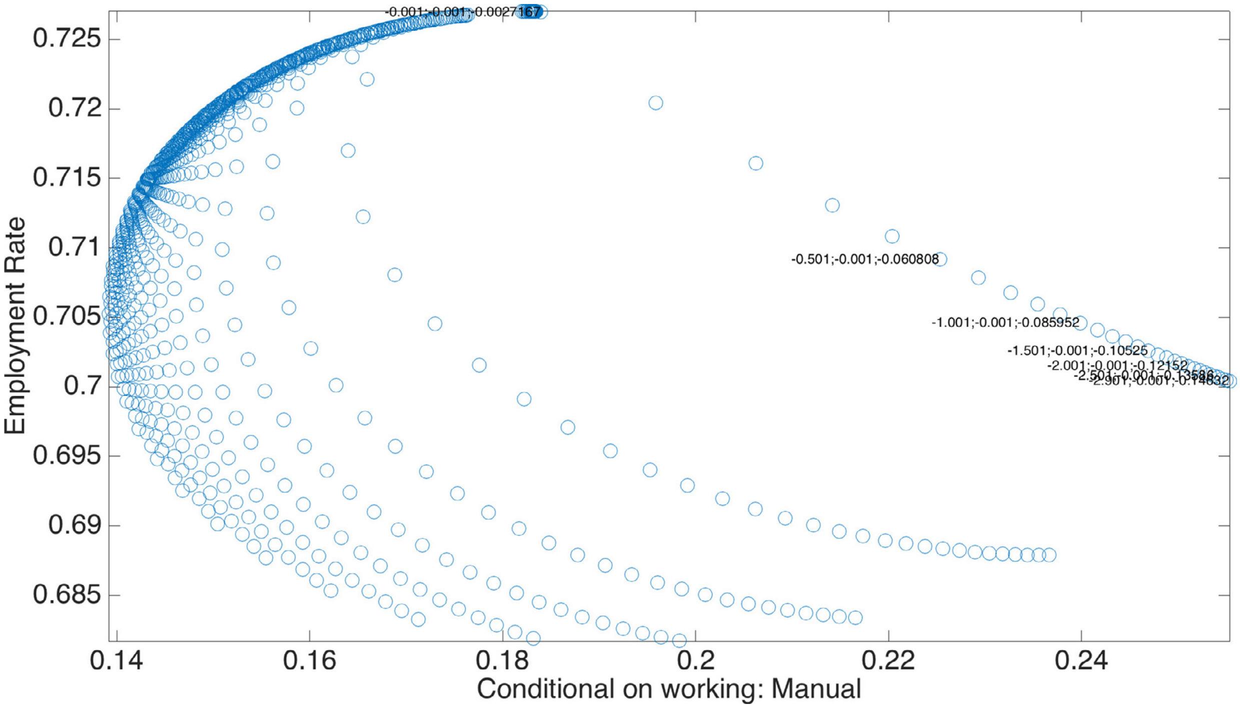

Fig. 4. Employment rate and occupation choice: substitutes. *Notes:* Each dot corresponds to different values of $(\eta_{G_R,R}, \eta_{G_{L_M},L_M}, \eta_{G_{L_M},R})$ as indicated to the right of selected dots. All simulations use a Pareto distribution for ability and $\mu = 1.3$. In 1989, the employment rate was 72.7%, and conditional on employment, the fraction working in manual occupations was 18.4%. In 2014 the employment rate fell to 64.9%, and conditional on employment, the fraction in manual occupations increased to 30.9%, off the scale of the figure in the "south–east" direction.

6.4. *What would it take?*

Finally, we ask what combination of parameter values and automation shock magnitude is required to account for *both* the observed occupational reallocation and employment change. With zero cross elasticities, doubling the *log deviation* shock (i.e. using $\widehat{A} = 2$) comes close to matching the empirical changes. This is depicted in Fig. A.4 in the Online Appendix. This is similarly true for the case when factors are complements or substitutes, results we make available upon request.

It is important to understand what these shock magnitudes mean. In our benchmark specification, $\widehat{A} = 1$ in log terms, implying that ICT capital has nearly tripled $(\exp(1) = 2.71)$ in *levels* relative to a balanced growth trend. A value of $\widehat{A} = 2$ implies a *greater than seven fold* $(\exp(2) = 7.38)$ increase. This highlights the challenge faced by advances in automation, as represented by measured changes in the stock of (quality adjusted) ICT capital, as the single force responsible for the changes in labor market outcomes experienced by the key demographic group studied here.

7. Conclusions

The share of employment in middle-skilled occupations has experienced a strong decline over recent decades. In this paper we show that this is primarily due to a fall in the propensity to work in these occupations conditional on demographic characteristics, rather than being driven by changes in the demographic composition of the economy. These propensity changes are concentrated among a relatively small subset of workers, who have experienced an increase in their propensity for non-employment (unemployment or non-participation) and their propensity to work in low-paying non-routine manual occupations. In fact, we show that these groups can account for a substantial fraction of the aggregate increase in non-employment and non-routine manual employment.

To shed light on the role of advances in automation technology in accounting for these phenomena, we study a flexible neoclassical model of the labor market, with endogenous occupation and participation decisions driven by worker heterogeneity. We show analytically that advances in automation cause workers to leave routine occupations and sort into non-employment and non-routine manual jobs. However, in the quantitative cases we consider, advances in automation technology on their own are unable to jointly generate changes in occupational shares and employment propensities that are quantitatively similar to those observed in the data for our key demographic group. We conclude that within the neoclassical context, accounting for a significant fraction of the changes along both margins requires relatively extreme combinations of parameter values and automation shock magnitude.

These results raise the question of what forces can account for our empirical findings. This paper has concentrated solely on the impact of automation. However, other changes have occurred in the U.S. economy that could have affected occupational choice and employment during the time period under study. Potentially relevant factors that we have abstracted from include changes in the share of high-skilled workers and their occupational choice, outsourcing and trade, and changes in policy affecting the incentive to participate in the labor market. In our view, the generality of our model provides a useful template for future quantitative research in evaluating the role of automation and other factors in contributing to the labor market outcomes discussed in the paper.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.jmoneco.2017.09.006.

References

Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: implications for employment and earnings. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, 4B. Elsevier, pp. 1043–1171.
Acemoglu, D., Autor, D., Dorn, D., Hanson, G.H., Price, B., 2016. Import competition and the great US employment sag of the 2000s. J. Labor Econ. 34 (S1), S141–S198.
Autor, D., Dorn, D., 2009. This job is “getting old”: measuring changes in job opportunities using occupational age structure. Am. Econ. Rev. 99 (2), 45–51.
Autor, D.H., Dorn, D., 2013. The growth of low skill service jobs and the polarization of the U.S. labor market. Am. Econ. Rev. 103 (5), 1553–1597.
Autor, D.H., Dorn, D., Hanson, G.H., 2015. Untangling trade and technology: evidence from local labour markets. Econ. J. 125 (584), 621–646.
Autor, D.H., Levy, F., Murnane, R.J., 2003. The skill content of recent technological change: an empirical exploration. Q. J. Econ. 118 (4), 1279–1333.
Beaudry, P., Green, D., Sand, B., 2016. The great reversal in the demand for skill and cognitive tasks. J. Labor Econ. 34 (S1), S199–S247.
Blinder, A., 1973. Wage discrimination: reduced form and structural estimates. J. Hum. Resour. 8, 436–455.
Charles, K.K., Hurst, E., Notowidigdo, M.J., 2013. Manufacturing Decline, Housing Booms, and Non-Employment. NBER Working Paper No. 18949. NBER.
Cortes, G.M., 2016. Where have the middle-wage workers gone? A study of polarization using panel data. J. Labor Econ. 34 (1), 63–105.
Cortes, G.M., Jaimovich, N., Nekarda, C.J., Siu, H.E., 2015. The micro and macro of disappearing routine jobs: a flows approach. University of British Columbia, Working paper.
Eden, M., Gaggl, P., 2016. On the Welfare Implications of Automation. University of North Carolina. Working Paper.
Flood, S., King, M., Ruggles, S., Warren, J.R., 2015. Integrated Public Use Microdata Series, Current Population Survey: Version 4.0. [Machine-Readable Database]. University of Minnesota, Minneapolis.
Gaggl, P., Wright, G.C., 2016. A short-run view of what computers do: evidence from a U.K. tax incentive. Am. Econ. J.: Appl. Econ 9 (3), 262–294.
Goos, M., Manning, A., 2007. Lousy and lovely jobs: the rising polarization of work in Britain. Rev. Econ. Stat. 89 (1), 118–133.
Goos, M., Manning, A., Salomons, A., 2009. Job polarization in Europe. Am. Econ. Rev.: Pap. Proc. 99 (2), 58–63.
Goos, M., Manning, A., Salomons, A., 2014. Explaining job polarization: routine-biased technological change and offshoring. Am. Econ. Rev. 104 (8), 2509–2526.
Graetz, G., Michaels, G., 2016. Robots at work. London School of Economics. Working paper.
Greenwood, J., Herkowitz, Z., Krusell, P., 1997. Long-run implications of investment-specific technological change. Am. Econ. Rev. 87, 342–362.
Gregory, T., Salomons, A., Zierahn, U., 2016. Racing With or Against the Machine? Evidence From Europe. ZEW – Centre for European Economic Research. Discussion Paper No. 16–053.
Hotz, V.J., Scholz, J.K., 2003. The earned income tax credit. In: Moffitt, R.A. (Ed.), Means-Tested Transfer Programs in the United States. University of Chicago Press, pp. 141–198.
Jaimovich, N., Siu, H.E., 2012. Job Polarization and Jobless Recoveries. NBER Working Paper 18334. NBER.
Krusell, P., Ohanian, L., Ríos-Rull, J.-V., Violante, G.L., 2000. Capital-skill complementarity and inequality: a macroeconomic analysis. Econometrica 68 (5), 1029–1053.
Lichter, A., Peichl, A., Siegloch, S., 2015. The own-wage elasticity of labor demand: a meta-regression analysis. Eur. Econ. Rev. 80, 94–119.
Mazzolari, F., Ragusa, G., 2013. Spillovers from high-skill consumption to low-skill labor markets. Rev. Econ. Stat. 95 (1), 74–86.
Michaels, G., Natraj, A., Van Reenen, J., 2014. Has ICT polarized skill demand? evidence from eleven countries over twenty-five years. Rev. Econ. Stat. 96 (1), 60–77.
Morin, M., 2016. Computer Adoption and the Changing Labor Market. Cambridge University. Working Paper.
Oaxaca, R., 1973. Male-female wage differentials in urban labor markets. Int. Econ. Rev. 14, 693–709.
vom Lehn, C., 2015. Labor Market Polarization, The Decline of Routine Work, and Technological Change: A Quantitative Analysis. Brigham Young University. Working Paper.