The Dynamics of Disappearing Routine Jobs: A Flows Approach^{*}

Guido Matias Cortes York University and RCEA gmcortes@yorku.ca

Christopher J. Nekarda Federal Reserve Board christopher.j.nekarda@frb.gov Nir Jaimovich University of Zurich and CEPR nir.jaimovich@uzh.ch

Henry E. Siu University of British Columbia and NBER hankman@mail.ubc.ca

This Draft: February 21, 2020

Abstract

We use matched individual-level CPS data to study the decline in middle-wage routine occupations during the last 40 years, and determine how the associated labor market flows have evolved. The decline in employment in these occupations can be primarily accounted for by changes in transition rates from non-participation and unemployment to routine employment. We study how these transition rates have changed since the mid-1970s, and find that changes are primarily due to the propensity of individuals to make such transitions, whereas relatively little is due to demographic changes. We also find that changes in the propensity to transition into routine occupations account for a substantial proportion of the rise in non-participation observed in the U.S. in recent decades.

^{*}A previous version of this paper was circulated under the title: "The Micro and Macro of Disappearing Routine Jobs: A Flows Approach." We thank David Dorn, Giuseppe Moscarini, Anne Polivka, two anonymous referees, the Editors, and numerous seminar participants at Yale, Rochester, Sherbrooke, Aachen, IZA, the CIREQ Labor Workshop on Unemployment, the Ifo Conference on Macroeconomics and Survey Data, and the CEA, CMSG, SOLE and TASKS Conferences for their helpful comments and suggestions. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. Cortes and Siu thank the Social Sciences and Humanities Research Council of Canada for support.

1 Introduction

In recent decades, labor markets in the United States and other developed countries have become increasingly polarized: the share of employment in middle-wage occupations has declined, while employment in both high- and low-wage jobs has increased. This "hollowing out" of the middle of the wage distribution has been linked to the declining share of employment in occupations with a high content of *routine tasks* – those activities that can be performed by following a well-defined set of procedures (see, for instance, Autor et al. (2006), Goos and Manning (2007), Goos et al. (2009) and Acemoglu and Autor (2011)). The declining employment in routine-intensive occupations has in turn been attributed to the fact that new technologies are particularly effective at performing these types of tasks (Autor et al., 2003).¹

In spite of the growing literature on polarization, relatively little is known about the individual-level patterns underlying the decline of routine employment. We use matched data from the monthly Current Population Survey (CPS) to analyze transitions into and out of employment in routine occupations. Our view is that characterizing the process by which routine employment is disappearing serves as an important guide in formalizing and evaluating theories of job polarization. It is equally important to the understanding of the changing labor market opportunities faced by different demographic groups, which is crucial in assessing policy implications; for example, the appropriate policy response would differ if the decline is accounted for by changes in the occupational choices of new labor market entrants than if it is due to increasing exit rates out of the labor force of prime-aged workers from routine employment.

Thus, our goal is to compare the quantitative importance of transition rates into routine occupations (*inflow rates*) relative to transitions rates out of routine occupations (*outflow rates*) in the observed decline in routine employment.² The set of inflow rates that we consider include transitions from labor force non-participation and unemployment to routine employment. The set of outflow rates that we consider are the reverse transitions, i.e. flows from routine employment to labor force non-participation and unemployment.³

We begin the analysis in Section 2 by describing how we use data from matched CPS files to construct nationally representative labor market flows at a monthly frequency from 1976 to 2018. We classify individuals in each month according to their labor market status (employed, unemployed, or not in the labor force) and their current or most recent occupational group (non-routine cognitive, routine cognitive, routine manual or non-routine manual, detailed below), and track their transitions over time.

In order to determine which changes in transition rates are key in accounting for the decline of routine employment, we perform a series of counterfactual exercises in Section 3. In these counterfactuals we replace different transition rates (for example, the entry rate from out of the labor force into routine manual employment), with their values observed prior to the onset of polarization. Applying a lawof-motion equation and using these counterfactual transition rates, we obtain counterfactual values of routine employment which inform us about how routine employment would have evolved had a specific transition rate not changed. By comparing these counterfactual values to those observed in the data, we can determine how much of the fall in routine employment would have been prevented if particular entry

¹See also Firpo et al. (2011), Goos et al. (2014), and the references therein regarding the role of outsourcing and offshoring in job polarization.

 $^{^{2}}$ Throughout the paper, we focus on the stock of employment in particular occupations as a share of the total working-age population, rather than as a share of total employment.

 $^{^{3}}$ Due to the CPS redesign in 1994, job to job transition rates are not measured in a consistent manner over our period of interest and thus their relative importance cannot be quantified. We discuss this issue in further detail below.

or exit rates had remained at their pre-polarization levels.

Before commenting on our results, it is useful to relate our counterfactual approach to other methods in the literature. First, we note that our approach is similar in spirit to the literature analyzing the role of job finding rates and job separation rates in accounting for the dynamics of unemployment over the business cycle (e.g. Darby et al., 1985; Hall, 2006; Elsby et al., 2009; Shimer, 2012). The key difference is that we analyze the secular decline in routine employment rather than the cyclical fluctuations in the unemployment rate. Our approach is also related to decomposition methods in economics (see e.g. Fortin et al., 2011), where the contribution of different channels to the observed change in an outcome variable of interest are quantified by counterfactually holding certain factors constant. Since in our particular setting counterfactuals need not add up to the total change (as is the case in decomposition approaches), we opt to use the term counterfactuals when referring to the experiments that we carry out, and use the term decomposition whenever we use more "common" decomposition methods in economics.⁴

Turning to results, our key finding is that inflow rates are much more important than outflow rates in accounting for the decline of routine employment. Specifically, had the inflow rates remained at their pre-polarization levels, at least 40 percent of the fall in routine employment would have been prevented. In contrast, had the outflow rates not changed, less than 5 percent of the fall in routine employment would have been prevented.

These results are performed using aggregate transition rates. However, changes in aggregate transition rates can be the result of both changes in the demographic composition of the economy (given that different demographic groups have different transition rates), and changes in transition rates for given demographic groups. As such, when we perform counterfactuals that hold aggregate transition rates constant, we effectively also remove the changes that are driven solely by demographic change. To address this shortcoming, we distinguish between the role of demographic change and changes in transition propensities within demographic groups in Section 4.

We begin with a series of Oaxaca-Blinder decompositions (Oaxaca, 1973; Blinder, 1973), and show that the vast majority of the decline in the aggregate inflow rates to routine occupations are driven by propensity changes conditional on demographic characteristics. Next, we re-visit our counterfactual exercises in greater detail, allowing for transition rates to evolve heterogeneously across demographic groups, and taking into account the changing relative importance of different demographic groups in the U.S. economy.

Our first exercise gauges how much of the change in routine employment is due to changes in the demographic structure of the economy. To do so we allow each demographic group's transition rates across the different labor market statuses to evolve as they did in the data. However, we hold the size of each demographic group constant at its 1976 level. This experiment results in a counterfactual series of aggregate labor market status that differs from the empirically observed one due solely to the lack of demographic change. By comparing the resulting counterfactual and observed routine employment changes, we can attribute a fraction of the decline to demographic change. In this experiment, we find

⁴It is worth noting that our approach shares many of the limitations of standard decomposition methods. In particular, it is a "partial equilibrium" approach that neglects potential endogenous responses of agents in the economy to the counterfactual transition rates we consider, and their impact on labor market stocks (employment, unemployment, and labor force participation). For example, it is possible that if the transition rates into routine employment had not fallen, observed demographic transitions like the increase in schooling would have been different. Such a change in demographic composition would likely have affected the distribution of the population across different labor market stocks. Similarly, it is possible that such counterfactual transitions would have altered the wages and returns to specific occupations, altering the labor market evolution. Our approach is silent on these general equilibrium, feedback, effects.

that demographic change alone can account for less than a quarter of the total long-run decline in routine manual employment, and essentially none of the decline in routine cognitive occupations.

In our second set of counterfactuals, we allow for demographic change to occur as it did in the data, but we hold group-specific inflow and outflow rates constant at their pre-polarization era levels. Once again, we find that inflow rates are quantitatively much more important than outflow rates. Moreover, we show that changes in the inflow rates to routine employment not only account for a substantial fraction of the decline in these occupations, but they also account for nearly three quarters of the rise in nonparticipation observed over the last two decades. This indicates that these falling inflow rates to routine employment have *not* been matched with increasing inflow rates to *non-routine* occupations; rather, they have resulted in increased propensities to remain out of employment.

We conclude our analysis by investigating which demographic groups are key in driving the changes that we have documented. Our quantitative finding is that three groups are salient in terms of their importance in driving the long-run dynamics of aggregate routine employment: males, the young, and those with intermediate levels of education.

As discussed above, relatively little attention has been paid to the labor market dynamics underlying the phenomenon of job polarization.⁵ Two recent papers are related to our analysis. Foote and Ryan (2014) analyze worker flows in the context of polarization, distinguishing between routine workers employed in different industries. Their paper differs from ours in that their primary goal is to study the cyclical properties of these flows, rather than their relationship with the long-term decline in routine employment. Smith (2013) describes the evolution of a number of flows into and out of routine employment and performs steady-state counterfactuals to analyze the importance of different transition rates in the decline of routine employment. As such, our papers share a number of findings, including the importance of job finding rate changes. At the same time, our analysis differs in a number of ways. First, our analysis allows us to determine not only the extent of transition rate change over time, but also how this is decomposed into composition and propensity change. Our detailed counterfactuals allow us to disentangle the role of demographic change and propensity change in accounting for the decline of routine employment. Second, while Smith (2013) focuses primarily on transitions between unemployment and employment, we analyze transitions into and out of the labor force, which we find to be key in accounting for routine employment dynamics. Finally, by using data from the 1970s, we are able to analyze how transition rates have changed relative to their levels prior to the onset of job polarization.

2 Data

We use monthly data from the Current Population Survey (CPS) spanning January 1976 to December 2018. The CPS is the primary source of labor force statistics for the U.S., and is sponsored jointly by the Census Bureau and the Bureau of Labor Statistics (BLS). We obtain the micro-data as made available by IPUMS (Flood et al., 2018). We restrict the sample to individuals aged 16 to 75.

Individuals can be matched across consecutive months due to the fact that the CPS is a rotating sample: households in the survey are interviewed for four consecutive months, then leave the sample for

⁵Evidence based on changes at the local labor market level, rather than on individual-level worker flows, is provided by Autor and Dorn (2009) and Autor et al. (2015). Cortes (2016) uses panel data to analyze the occupational mobility patterns of workers switching out of routine jobs. Cortes et al. (2017) analyze how the propensity to work in routine occupations has changed for different demographic groups, but do not consider the worker flows underlying these changes.

eight months, before returning for another four months. Given this structure, up to 75% of sampled households are potentially matched in any given month. In practice, the fraction of households matched is slightly lower due primarily to the fact that the CPS is an address-based survey: households that move are not followed. Also, in certain months the CPS made changes to household identifiers, making it impossible to match individuals across these interviews.⁶ We match individuals across consecutive months based on their person identifier, making sure that they have consistent information for sex, race and age.

The main advantage of the CPS is its large sample size, designed to be representative of the U.S. population, allowing for the observation of individual-level transitions across labor market states at a monthly frequency. Another important advantage is its time coverage, spanning periods both before and after the onset of job polarization, as discussed below.⁷ The primary challenge of the CPS is the 1994 survey redesign that induced certain data discontinuities, which we discuss below. The remainder of this section describes how we use the data to classify individuals according to their occupation and labor force status, and how we construct transition rates across labor market states.

2.1 Labor Force and Occupation Categories

We categorize all individuals in the sample according to their labor force status: employed, unemployed, or not in the labor force. The CPS records employed workers' description of their current occupation in their main job, and also unemployed workers' description of their occupation in their most recent job (if they have ever worked before). The individual's description is then assigned a 3-digit occupation code.⁸ While occupational data exists for the employed and unemployed, this is not the case for those classified as out of the labor force.⁹ We are therefore constrained to consider only one labor force non-participation category that does not distinguish based on previous occupation.

Following the recent literature (e.g., Acemoglu and Autor (2011), Autor and Dorn (2013), Cortes (2016), Jaimovich and Siu (2012)), we consider four broad occupational groups. We do this by delineating occupations along two dimensions of the characteristics of tasks performed on the job: "cognitive" versus "manual," and "routine" versus "non-routine." The distinction between cognitive and manual occupations is straightforward, based on differences in the extent of mental versus physical activity. The distinction between routine and non-routine jobs 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 and procedures, the occupation is considered routine. If instead the job involves a variety of tasks, requiring flexibility, problem-solving, or human interaction skills, the occupation is non-routine. As such, the four occupational groupings are: non-routine cognitive, routine cognitive, routine manual and non-routine manual.

We aggregate detailed occupational codes into these four clusters based on broad occupational groupings. Specifically, *non-routine cognitive* occupations are Professional, Managerial and Technical Occupa-

⁶This occurs in January 1978, July 1985, October 1985, January 1994, June 1995 and September 1995.

⁷By contrast, while the Panel Study of Income Dynamics tracks individuals over a longer time period, its sample is much smaller (making it problematic for the analysis of occupational and demographic detail) and available only at an annual or bi-annual frequency. The Survey of Income and Program Participation is at a monthly frequency and has, in certain waves, sample sizes comparable to the CPS; however, it begins after the onset of job polarization.

⁸For matched individuals who are unemployed and have a missing occupation code, we impute their previous month's occupation code, if it is available. We make the imputation for several consecutive months, if necessary.

⁹The exception is when they are in the 'outgoing rotation group' (i.e., in their fourth or eight month in the sample) but this information is not useful as we cannot match these individuals to the following month.

tions; routine cognitive are Sales and Clerical Occupations; routine manual are Production, Craft and Repair Occupations, Operators, and Transportation and Material Moving Occupations; and non-routine manual are Service Occupations. The occupation codes changed in 1983, 1992, 2003 and 2011, when the CPS moved between the 1970, 1980, 1990, 2000, and 2010 classification systems. To maintain consistency through time, we map each 3-digit occupation code across the five classification systems used by the BLS since 1976 into the four occupation categories; details of the mapping are in Appendix Table A.1.¹⁰ Throughout the paper, we exclude observations for those working in the military, and those in farming, fishing, and forestry occupations.

Given our occupational groups, we can classify individuals into one of nine mutually exclusive categories: employed in one of the four occupation groups (denoted ENRC, ERC, ERM, and ENRM for non-routine cognitive, routine cognitive, routine manual, and non-routine manual occupations, respectively); unemployed with previous job in one of the four occupation groups (UNRC, URC, URM, and UNRM); or not in the labor force (NLF).¹¹

Table 1 presents descriptive statistics for the full sample. Panel A displays information for the period before 1990 (1976–1989), while Panel B displays information for the more recent period (1990–2018). As is evident, there is clear heterogeneity across occupations in terms of demographic composition. For instance, the level of educational attainment is highest in non-routine cognitive occupations, and lowest in non-routine manual ones; routine occupations tend to employ middle-skilled workers (high school graduates and those with some college education). Similarly, there is clear heterogeneity in gender composition: while workers in routine cognitive occupations are predominantly female, routine manual ones are predominantly male.

Figure 1 displays the monthly time series of employment in each occupational group as a share of the total working-age population. Despite our effort to define groups consistently, there is an obvious discontinuity in 1983 with the introduction of the 1980 occupation codes.¹² The discontinuity re-allocates employment from non-routine cognitive occupations to routine cognitive. In spite of this, the figure illustrates the obvious rise in non-routine employment (both cognitive and manual).

The dynamics of routine manual and routine cognitive employment are quite different. Employment in routine manual occupations (ERM) begins to disappear in the early 1980s. The business cycle dimension discussed in Jaimovich and Siu (2012) is clearly evident: employment in these occupations falls during the back-to-back recessions of 1980/82, the recessions of 1991 and 2001, and the recent Great Recession, and fails to recover during the subsequent expansions. By contrast, employment in routine cognitive occupations (ERC) grows through the 1980s, before reversing in the early 1990s.¹³ Its decline and lack

¹⁰We have also categorized occupations using the crosswalk of Autor and Dorn (2013), itself an adaptation of Meyer and Osborne (2005). This methodology converts all of the 3-digit occupation codes from the 1970, 1980, 1990, and 2000 systems to a common coding system (we developed our own crosswalk to convert the 2010 codes). The common codes are then aggregated into the four broad categories. The results from that methodology are largely similar relative to those from the current mapping; we refer the reader to the November 2013 version of our paper for details. However, using the Autor and Dorn (2013) crosswalk generates noticeable discontinuities in the non-routine cognitive and routine cognitive groups between the 1990 and 2000 classification systems; these discontinuities are avoided by the current methodology.

 $^{^{11}}$ There is a small group of individuals who are unemployed with no previous occupational information. For simplicity we remove them from the analysis.

 $^{^{12}}$ Because of the major changes instituted between the 1970 and 1980 classification systems, this discontinuity is a feature of all categorization methodologies that assign 3-digit level codes to one of the occupation groups. See, for instance, the discussion of the Autor and Dorn (2013) methodology in the November 2013 version of this paper, and also Jaimovich and Siu (2012) for further discussion.

 $^{^{13}}$ Note that this pattern is not sensitive to the occupational grouping method used in this paper. For example, Figure 13a in Acemoglu and Autor (2011) shows an increase in the share of employment in routine cognitive occupations (clerical

	Panel A: 1976-1989						
	Full	ENRC	ERC	ERM	ENRM	NLF	
Average age	40.20	39.71	36.53	36.63	35.43	46.72	
Fractions within the	occupation g	roup					
HS dropouts	0.28	0.06	0.10	0.33	0.36	0.44	
HS graduates	0.56	0.43	0.75	0.63	0.59	0.48	
College graduates	0.16	0.52	0.15	0.04	0.05	0.09	
Male	0.48	0.60	0.32	0.82	0.42	0.30	
Non-White	0.13	0.09	0.10	0.13	0.20	0.15	
Married	0.64	0.72	0.62	0.68	0.52	0.62	
Total number of obs	ervations (ma	illions)					
Unweighted	17.60	3.25	2.92	3.25	1.60	5.85	
Share of sample							
Weighted	1.00	0.18	0.17	0.19	0.09	0.33	

		Panel B	: 1990-2018					
	Full	ENRC	ERC	ERM	ENRM	NLF		
Average age	42.04	42.49	39.44	39.73	37.25	46.98		
Fractions within the	e occupation g	roup						
HS dropouts	0.16	0.02	0.06	0.19	0.20	0.28		
HS graduates	0.59	0.37	0.72	0.75	0.70	0.57		
College graduates	0.25	0.61	0.22	0.07	0.10	0.16		
Male	0.49	0.50	0.36	0.84	0.43	0.38		
Non-White	0.19	0.16	0.17	0.17	0.24	0.21		
Married	0.57	0.68	0.56	0.61	0.48	0.53		
Total number of observations (millions)								
Unweighted	33.60	7.92	5.52	4.98	3.49	10.20		
Share of sample								
Weighted	1.00	0.23	0.16	0.15	0.11	0.31		

Note: The full sample excludes military workers, workers in farming, fishing, and forestry occupations, and unemployed individuals with unknown previous occupation. ENRC, ERC, ERM, and ENRM stand for employment in non-routine cognitive, routine cognitive, routine manual, and non-routine manual occupations; NLF stands for not in the labor force. HS graduates include those with some post-secondary, but less than college degree. The four unemployment categories (not displayed) account for the remaining 4.2% of the sample in the 1976–1989 period, and 3.8% of the sample in the 1990–2018 period.

of recovery are evident following the 1991 and 2001 recessions. This pattern is repeated in a dramatic manner beginning in 2007: a sharp disappearance in the Great Recession with no recovery since. Our analysis focuses on the transition patterns underlying the decline in these two categories of routine employment, taking into account the differences in timing.

2.2 Construction of Transition Rates

Using the individual-level information on labor force status and occupation, we construct monthly transition rates across the nine labor market states. The date t transition rate between labor market state A and state B is computed as the number of individuals switching from A at date t to B at date t + 1,

and sales) up until 1989 before reversing thereafter.



Figure 1: Employment Stocks in Monthly CPS Data

Note: Each employment stock is measured as a fraction of the working-age population. The mapping of 3-digit occupation codes to the four broad categories is detailed in Appendix Table A.1.

divided by the number of individuals in state A matched between dates t and t + 1.¹⁴ This generates a 9×9 matrix of transition rates, ρ_t , for each month t in our sample. This matrix can be split into sub matrices as follows:

$$\rho_{t} = \begin{bmatrix}
\rho_{t}^{EE} & \rho_{t}^{UE} & \rho_{t}^{NE} \\
\rho_{t}^{EU} & \rho_{t}^{UU} & \rho_{t}^{NU} \\
\rho_{t}^{EN} & \rho_{t}^{UN} & \rho_{t}^{NN}
\end{bmatrix},$$
(1)

where

- ρ_t^{EE} (4×4): employment "stayer" rates and "job-to-job" transition rates across occupation groups;
- ρ_t^{EU} (4 × 4): transition rates from employment to unemployment, or "job separation rates;"

¹⁴That is, individuals who leave the sample between t and t+1 (outgoing rotation group, attritioners) are excluded from the computation of transition rates. Matched individuals are weighted using CPS sample weights from month t.

- ρ_t^{EN} (1 × 4): transition rates from employment to non-participation;
- ρ_t^{UE} (4 × 4): transition rates from unemployment to employment, or "job finding rates;"
- ρ_t^{UU} (4 × 4): unemployment stayer rates;
- ρ_t^{UN} (1 × 4): transition rates from unemployment to non-participation;
- ρ_t^{NE} (4 × 1): transition rates from non-participation to employment,
- ρ_t^{NU} (4 × 1): transition rates from non-participation to unemployment,
- ρ_t^{NN} (1 × 1): non-participation stayer rates.

The law-of-motion governing the evolution of the labor market "stocks" is given by:

$$\underbrace{Stocks_{t+1}}_{(9\times1)} = \underbrace{\rho_t}_{(9\times9)} * \underbrace{Stocks_t}_{(9\times1)}$$
(2)

where $Stocks_t = [ENRC_t \ ERC_t \ ERM_t \ ENRM_t \ UNRC_t \ URC_t \ URM_t \ UNRM_t \ NLF_t]'$ is the vector summarizing the fraction of working age population in each state. To understand the dynamics implied by Equation (2), consider the evolution of employment in routine manual occupations. The monthly net change from t to t + 1 depends on the "inflows" of individuals into ERM from unemployment, non-participation, and employment in other occupations, relative to the "outflows" from ERM to unemployment, non-participation, and employment in other occupations. Equation (2) summarizes these flows by the size of each of the stocks and the corresponding transition rates between them.

We determine which transitions in ρ_t are particularly important in accounting for the decline in routine employment, by performing a number of counterfactual exercises discussed in the next section. Before proceeding, we illustrate that the law-of-motion provides a good approximation of the stocks measured cross-sectionally, as presented in Figure 1. This may not be the case as Equation (2) relies only on an initial measure of stocks and iterates forward using the subsequent transition rates. Transition rates computed from matched individuals may fare poorly due to entry and exit from the sample (attrition and rotation of sampled individuals).

Figure 2 plots employment in each occupation group and labor force non-participation from 1976:1 to 2018:12. The stocks based on the full sample are the blue, solid lines; the estimates based on the law of motion in Equation (2) are the red, hatched lines.¹⁵ The time series derived from transition rates slightly underestimate the fraction of employed workers, and overestimate the fraction out of the labor force.¹⁶ By the end of our sample period, the labor force non-participation rate is overestimated by approximately three percentage points. Interestingly, the gap in employment is due entirely to an underestimation of the fraction of the labor force non-participation, and this discrepancy has become more pronounced in the last three years of our data. Nonetheless, the stocks based on the law-of-motion follow the long-run

 $^{^{15}}$ For months where we have missing data because of the change in CPS sample identifiers or occupational coding systems, we keep the transition rates as missing, leaving stocks constant.

 $^{^{16}}$ We note that "margin error," as documented by Frazis et al. (2005) and others, generates a qualitatively similar discrepancy when stocks are constructed by adding and subtracting gross flows in matched CPS data. Margin error discrepancies accumulate over time. By contrast, discrepancies from our procedure do not accumulate. Relative to the cumulative addition/subtraction of gross flows, we iterate on transition rates defined only for individuals who are matched across consecutive months. As such, our procedure essentially imputes to those who leave the sample the same transition probabilities as those who remain.



Figure 2: Labor Market Stocks from Full Sample and based on Law of Motion

Note: Each stock is measured as a fraction of the working-age population. The mapping of 3-digit occupation codes to the four broad categories is detailed in Appendix Table A.1. The series based on the law of motion uses transition rates obtained from matched monthly CPS samples and the law of motion in Equation (2).

paths of those based on the full data quite closely, and particularly so in the case of routine employment. This rationalizes our approach of focusing on transition rates derived from labor market flows in order to understand the long-run disappearance of routine jobs.

The main data challenge that arises when analyzing the importance of different labor market flows using matched CPS data is the discontinuity induced by the CPS survey redesign. In 1994, the CPS switched to a method of dependent interviewing to ease respondent burden and improve data quality. For occupation data, interviewers asked whether the interviewee had the same job as in the previous month; if the answer was yes, the individual would automatically receive the same occupation code. Dependent coding substantially reduced the occurrence of spurious transitions across occupations at the monthly frequency (see Kambourov and Manovskii (2013) and Moscarini and Thomsson (2007)). This generates a discrete break in the measured transition rates across occupations for those reporting employment in consecutive months.¹⁷ This break limits our ability to study the role of changes in job-to-job flows between occupations over time. We therefore restrict our attention in the remainder of the paper to the flows to and from routine employment that are measured in a consistent manner throughout our sample period, namely the inflows from unemployment and non-participation into routine employment, and the outflows from routine employment to unemployment and non-participation.

3 The role of inflows and outflows to routine employment

3.1 Aggregate Counterfactual

In this section, we study the role of changes in aggregate transition rates into and out of routine employment in accounting for its decline over the past 40 years. Rather than simply looking at the change in transition rates over time, we use counterfactual analysis since the evolution of routine employment depends on the volume of inflows and outflows to and from this labor market state, which are themselves a product of the transition rates *and* the stocks of all the various labor market states. Thus, a relatively large change in a transition rate might have little quantitative effect on routine employment if the transition rate is small to begin with, or if the source stock is small (e.g., one of the unemployment categories). On the other hand, a small change in a transition rate could have a substantial impact if the source group is large (e.g., labor force non-participants). Our approach accounts for these considerations.

As a first step, given that transition rates vary significantly over the business cycle, we divide the sample period into recessionary phases (based on NBER peak to trough dates) and non-recessionary phases (which include all other months in the sample). Table 2 lists the 11 individual phases from 1976 to 2018. We denote the six expansions as E1 through E6, and the five recessions as R1 through R5. We then calculate the average of each transition rate within each phase. In our analysis, we replace the average value of specific transition rates during the post-polarization period with the average from the corresponding pre-job polarization phases (which we discuss in detail below).

Figure 3 plots the time series of the two routine employment stocks when using the average transition rates for each of the 11 phases in the law of motion from Equation (2).¹⁸ We call these the *stocks based*

 $^{^{17}}$ The 1994 redesign also induces a discontinuity in the measured transition rates between non-participation and unemployment.

¹⁸ Due to the January 1994 redesign of the CPS and the discontinuities that this induces in certain transition rates, the averages for phase E4 used in this section are calculated over the period 1994:1 to 2001:2.

	Table 2:	List	of	individual	business	cvcle	phases
--	----------	------	----	------------	----------	-------	--------

Expansions:	Recessions:
1976m1-1979m12 (E1)	$1980 \mathrm{m1}$ - $1980 \mathrm{m7}$ (R1)
1980m8-1981m6 (E2)	1981m7-1982m11 (R2)
1982m12-1990m6 (E3)	$1990 \mathrm{m7}$ - $1991 \mathrm{m3}$ (R3)
1991m4-2001m2 (E4)	2001m3-2001m11 (R4)
2001m12-2007m11 (E5)	2007m12-2009m6 (R5)
2009m7-2018m12 (E6)	

Note: The phase numbers as referred to throughout the text are given in parentheses.

on average rates. The figure also plots the stocks based on the monthly rates as shown in Figure 2, as well as a band-pass filtered version of this series.¹⁹ The series based on the monthly rates and the one based on phase averages differ to the extent that the average transition rates abstract from fluctuations within a phase. Evidently, the stocks based on average rates provide a good approximation of the data; eliminating the high frequency movements does not obscure the dynamics underlying the long run decline in routine employment.

Our goal is to understand how much of the declines in routine employment can be accounted for by changes in different transition rates to and from these occupations. To do this, we perform a series of counterfactual experiments allowing all average transition rates to evolve as observed in the data *except* for certain ones that are held constant at their pre-polarization averages. In choosing the phases representative of the pre-polarization era, we account for the difference in timing of when routine cognitive and routine manual employment begin to decline. The expansion of the late-1970s and the recession in 1980 are set as the pre-polarization phases for routine manual occupations. Hence, in the counterfactual exercises where we hold transition rates to and from *ERM* fixed, we replace: (i) their average value during recessions R2 through R5 with the average for R1, and (ii) their average value during expansions E2 through E6 with the average for E1. The decline of routine cognitive employment (*ERC*) occurs later on. Hence, we choose R3 and E3 as the benchmark pre-polarization periods of the transition rates. The results are essentially unchanged if we allow the transition rates during recessionary periods, which tend to be short lived, to evolve as in the data; that is, the key to the analysis is whether the transition rates during expansionary periods change or not.

To illustrate the counterfactual exercise more formally, consider the law-of-motion in Equation (2): $Stocks_{t+1} = \rho_t * Stocks_t$. In the counterfactual experiments, we compute counterfactual series of stocks given by:

$$Stocks_{t+1}^{CF} = \rho_t^{CF} * Stocks_t^{CF}$$
(3)

where ρ_t^{CF} is a counterfactual matrix of transition rates, and where in the initial period we use the observed stocks in the data: $Stocks_0^{CF} = Stocks_0$.

For illustrative purposes, suppose that we are interested in the role of the inflow rates from nonparticipation to routine manual employment. Recall that ρ_t is a (9 x 9) matrix with the transition rates from each source labor market state to each destination state (see Equation (1)). In the counterfac-

 $^{^{19}}$ Our band pass filter removes fluctuations at frequencies higher than 18 months (business cycle fluctuations are traditionally defined as those between frequencies of 18 and 96 months). We implement this using the band pass filter of Christiano and Fitzgerald (2003), who discuss the merits of their method for isolating fluctuations outside the traditional business cycle frequencies and near the endpoints of datasets.

Figure 3: Routine Employment Stocks based on Law of Motion using Monthly Rates and Phase Averages



Note: Each employment stock is measured as a fraction of the working-age population. The mapping of 3-digit occupation codes to the two routine employment categories is detailed in Appendix Table A.1.

tual where we explore the role of inflows from non-participation (NLF) to routine manual employment (ERM), all elements of the matrix ρ_t^{CF} are the same as in ρ_t , except that we set

$$\rho_t^{NLF,ERM} = \rho_{CF}^{NLF,ERM},$$

i.e. we replace the transition rate from NLF to ERM with a counterfactual, time-invariant rate, which is equal to its value in the pre-polarization phase. The sum of the transition rates out of NLF would not sum to 1. We therefore allocate the difference between the observed and the counterfactual rate proportionally to all other transition rates out of the source labor market state (in this example NLF) according to their relative magnitude. That is, for example, the transition rate to ENRC, $\rho^{NLF,ENRC}$, in the counterfactual transition matrix ρ_t^{CF} would become

$$\rho_{t,CF}^{NLF,ENRC} = \rho_t^{NLF,ENRC} + \frac{\rho_t^{NLF,ENRC}}{1 - \rho_t^{NLF,ERM}} \times \left(\rho_t^{NLF,ERM} - \rho_{CF}^{NLF,ERM}\right).$$

All other transition rates out of NLF are adjusted accordingly.

Once we have constructed the relevant counterfactual matrix ρ_t^{CF} for the experiment under consideration, counterfactual stocks $Stocks_{t+1}^{CF}$ are obtained using Equation (3). We then compute the total

counterfactual change in the labor market state of interest (in this example ERM) from its peak to its level at the end of the sample (which coincides with its minimum value in our data), period T:

$$\Delta ERM^{CF} = ERM_T^{CF} - ERM_{Peak},$$

and compare this to the actual change observed in the data:

$$\Delta ERM = ERM_T - ERM_{Peak}.$$

We summarize the role of the transition rate under consideration (in this example the inflow rate from NLF to ERM) in accounting for the decline in the labor market stock of interest (in this example ERM) by computing the fraction of the decline avoided in the counterfactual experiment, i.e.

$$1 - \frac{\Delta ERM^{CF}}{\Delta ERM}.$$
 (4)

All of the counterfactual experiments that we report below are variations of this. Each experiment considers a different counterfactual matrix ρ_t^{CF} , depending on the transition rate being considered, and on the labor market stock of interest (either *ERM* or *ERC*).

Overall, the stocks based on average rates exhibit a fall in ERM from 19.9 percent to 12.4 percent of total employment, and a fall in ERC from 18.3 percent to 14.2 percent. Our counterfactual exercises ask how much of these 7.5 and 4.1 percentage points, respectively, would have been avoided if certain transition rates had remained at their pre-polarization levels.

3.2 Results

The key results are presented in Table 3. Our first experiment sets all of our inflow and outflow rates of interest to their pre-job polarization levels. This holds the following constant: (i) the inflow rates from all categories of unemployment into employment in a routine occupation, (ii) the inflow rates from non-participation into employment in a routine occupation, (iii) the outflow rates from routine employment into unemployment, and (iv) the outflow rates from routine employment into non-participation. All other transition rates are allowed to evolve as they do in the data.

The table shows that holding these transition rates constant mitigates 40% of the decline in ERM employment from its peak in the early 1980s to its trough at the end of our sample. Meanwhile, 46% of the decline in ERC employment is mitigated in this experiment. Hence, 40-46% of the decline in routine employment can be accounted for through changes in the flows between routine employment and non-employment. The remaining 54-60% of the decline would be accounted for by changes in other transition rates (job to job changes, transitions to and from non-routine employment, and transitions between unemployment and non-participation). To assess the statistical significance of these results, we obtain a 90% confidence interval by repeating the experiment on a set of 100 bootstrapped samples. These confidence intervals are reported in square brackets below the respective estimates, and show that the results are strongly statistically significant.²⁰

Having established the joint role of inflows and outflows between routine employment and non-

 $^{^{20}}$ We thank the editor and an anonymous referee for suggesting this approach to us.

	Fraction of decline avoided		
	Employed:	Employed:	
	Routine Manual (ERM)	Routine Cognitive (ERC)	
Inflow and Outflow Rates	$\begin{array}{c} 0.40 \\ [0.37, 0.41] \end{array}$	$\begin{array}{c} 0.46 \\ [0.44, 0.49] \end{array}$	
Inflow Rates	$\begin{array}{c} 0.34 \\ [0.32,0.36] \end{array}$	$0.43 \\ [0.41, 0.46]$	
Inflow from Non-Participation	$\begin{array}{c} 0.19 \\ [0.17, \ 0.21] \end{array}$	$\begin{array}{c} 0.29 \\ [0.27, \ 0.31] \end{array}$	
Inflow from Unemployment	$\begin{array}{c} 0.16 \\ [0.14, 0.18] \end{array}$	$\begin{array}{c} 0.16 \\ [0.13, \ 0.19] \end{array}$	
Outflow rates	$0.04 \\ [0.03, \ 0.06]$	$\begin{array}{c} 0.04 \\ [0.03, 0.05] \end{array}$	
Outflow to Non-Participation	0.05 [0.03, 0.08]	$\begin{array}{c} 0.06 \\ [0.05,\ 0.08] \end{array}$	
Outflow to Unemployment	-0.01 [-0.03, 0.02]	-0.02 [-0.03, 0.00]	

Table 3: Results from aggregate counterfactual exercises

Note: The table reports the fraction of the decline in ERM and ERC (from their respective peaks to their levels in 2018:12) that is avoided by holding different sets of transition rates constant at their pre-polarization levels. The ERM and ERC stocks are measured as a fraction of the working-age population. The numbers in square brackets represent the 90% confidence interval obtained from 100 bootstrap simulations of the counterfactual experiments.

employment, we next assess the relative importance of different subsets of these transition rates. As Table 3 shows, the majority of the overall effect that we find can be attributed to changes in inflow rates, while the role of outflow rates is small. By fixing the inflow rates to routine employment, 43% of the fall in *ERC* would have been prevented (vs 46% when both inflow and outflow rates are considered). In the case of *ERM*, inflow rates prevent 34% of the fall (vs 40% when both inflow and outflow rates are considered). By contrast, holding the outflow rates from routine employment constant mitigates only 4% of the fall in routine employment.

The table further breaks down the role of inflows from unemployment relative to inflows from nonparticipation. In the case of ERM we find that both entry margins are of roughly equal importance in accounting for the joint 34% effect reported above. For ERC, by contrast, changes in the inflow rate from non-participation are much more important relative to that from unemployment. The breakdown for outflows confirms that neither outflows to unemployment nor non-participation play a major role in the decline of routine employment. In fact, changes in the outflow rate to unemployment provide a negative contribution—if the outflow rates to unemployment had remained at their pre-polarization levels, the decline of routine employment would have been stronger than observed in the data.

To summarize, the key takeaway is that the decline in routine employment is primarily due to a reduction in the inflow rate to these jobs among the unemployed and those out of the labor force. Separations from routine employment towards non-employment play little to no role in driving the observed decline in these occupations.

4 Demographic Breakdown

The previous section finds that changes in inflow rates account for a substantial fraction of the disappearance of employment in routine occupations. This is based on counterfactual exercises where specific *aggregate* transition rates are held constant at their pre-polarization levels. It is well known that labor market transition rates vary significantly across demographic groups. For instance, young individuals are more likely to transit from unemployment to employment relative to those who are older. Changes in the demographic composition of the U.S. economy could potentially account for some of the changes in routine employment transition rates observed over the past 40 years.

In this section we disentangle the extent to which the key changes identified in Section 3 can be attributed to changes in: (i) the demographic composition of the U.S. economy, and (ii) the propensities to make certain transitions for individuals from particular demographic groups. If transition rate changes were due principally to demographic shifts, such as the aging of the U.S. population, one might argue that polarization is a natural consequence of demographic change.²¹ By contrast, if the changes are due principally to changes in propensities and vary across routine and non-routine occupations, it suggests attribution to forces responsible for job polarization.

We proceed as follows. First, we perform a series of Oaxaca-Blinder decompositions (Oaxaca, 1973; Blinder, 1973) and decompose changes in transition rates into a component that can be explained by changes in the demographic composition of different labor market states, and a component related to changes in transition propensities, conditional on demographic characteristics. We then return to counterfactual exercises to determine the importance of demographic change, and the importance of changes in transition propensities among particular demographic groups, in the decline of routine employment.

4.1 Oaxaca-Blinder Decompositions of Aggregate Inflow Rates

Let ρ_{it}^{AB} be a dummy variable defined for all individuals who are in labor market state A in period t. It is equal to 1 if individual *i* switches from state A to state B between month t and month t + 1, and is equal to zero otherwise. The average aggregate transition rate for a given expansionary phase τ corresponds to the average of ρ_{it}^{AB} , which we denote $\bar{\rho}_{\tau}^{AB}$. The transition rate probability ρ_{it}^{AB} can be specified as a function of demographic characteristics as follows:

$$\rho_{it}^{AB} = X_{it}^A \beta_\tau + \epsilon_{it},\tag{5}$$

where X_{it}^A comprises a set of standard demographic variables available in the CPS, and β_{τ} represents a set of phase-specific coefficients. Estimating this linear probability model for each expansionary phase allows us to obtain phase-specific estimates of β_{τ} , which we can use to perform a standard Oaxaca-Blinder (OB) decomposition of the change in the aggregate transition rate over time:

 $^{^{21}}$ Such an argument is valid for demographic composition changes that are orthogonal to changes in the labor market. The argument is less clear cut along other dimensions; for instance, it could be argued that rising educational attainment has been driven to some extent by the rise of non-routine cognitive vs routine manual job opportunities. Such issues cannot be settled simply within this empirical framework.

$$\overline{\rho}_{0}^{AB} - \overline{\rho}_{1}^{AB} = \left(\overline{X}_{0}^{A}\widehat{\beta}_{0}\right) - \left(\overline{X}_{1}^{A}\widehat{\beta}_{1}\right)$$
$$= \left(\overline{X}_{0}^{A} - \overline{X}_{1}^{A}\right)\widehat{\beta}_{0} + \left(\overline{X}_{1}^{A}\right)\left(\widehat{\beta}_{0} - \widehat{\beta}_{1}\right).$$
(6)

The change in the transition rate across phases 0 and 1 (on the left-hand side of the equation) can be decomposed into two parts. The first, given by the first term in Equation (6), is the component attributed to changes in the demographic composition of individuals in labor market state A across phases 0 and 1. The second part is attributed to changes in $\hat{\beta}_{\tau}$, reflecting changes in the propensities to transition from state A to B for particular demographic groups. We thus decompose transition rate changes from the pre- to the post-polarization era into changes that are "explained" or "unexplained" by observables.

We focus on two key transition rates we have identified in driving the decline in routine employment. The first is the rate of inflows from non-participation. The second is the "return" job finding rate, the probability that unemployed workers previously working in routine cognitive (routine manual) occupations return to routine cognitive (routine manual) employment; because the stock of unemployed is relatively small and because the incidence of occupational group switching is relatively rare, this is the quantitatively relevant inflow rate from unemployment. Our vector of demographic characteristics X_{it}^A includes controls for age (six age bins: 16-24, 25-34, ..., 55-64, and 65+), education (less than high school, high school diploma or some post-secondary, and college graduate), gender, and race (white versus other). Given that the CPS data are not seasonally adjusted, we also include controls for seasonality (a full set of calendar month dummies).

Table 4 presents the decomposition results for the key inflow rates to routine manual employment. Panel A shows that in the baseline, pre-polarization period (the late 1970s expansion), the "return" job finding rate for routine manual workers was above 23%. This rate declines during the subsequent expansions, and in particular in the post-Great Recession period. Importantly, the decomposition shows that these declines are not driven by changes in the demographic characteristics of unemployed routine manual workers. In fact, changes in demographics would have predicted a slight increase in the return job finding rate (as indicated by the positive sign for the 'Composition' component). The decline is instead entirely driven by a fall in the propensity to return to routine manual employment, conditional on demographic characteristics (as indicated by the sign and the magnitude of the 'Propensities' component).

Panel B presents the analogous results for the inflow rate from non-participation. This rate was 1.23% in the pre-polarization period, and declines during all subsequent expansions. Although the declines are small (between 0.1 and 0.4 p.p.), they have an important impact given the large size of the non-participant pool. The decomposition results once again indicate that it is a fall in the propensity to make this labor market transition, rather than a change in the demographic characteristics of non-participants, that is responsible for the decline.

Table 5 presents the decomposition results for the inflow rates to routine cognitive employment. Panel A shows that the return job finding rate for unemployed routine cognitive workers falls from its pre-polarization benchmark level in all subsequent expansions, particularly after the Great Recession. In this case, both changes in the demographic composition of unemployed routine cognitive workers, and changes in propensities conditional on demographic characteristics, contribute to the decline. However, the contribution of propensity changes is much larger.

	Baseline Expansion (1976m1-1979m12): 23.30%				
-	1982m12-	1991m4-	2001m12-	2009m7-	
	1990m6	2001m2	2007m11	2018m12	
Total Change	-2.60^{***}	-1.22^{***}	-1.83^{***}	-7.34^{***}	
	(0.28)	(0.29)	(0.32)	(0.28)	
Composition	$+0.45^{***}$	$+1.02^{***}$	$+0.79^{***}$	+0.19	
	(0.07)	(0.10)	(0.15)	(0.14)	
Propensities	-3.05^{***}	-2.242^{***}	-2.61^{***}	-7.54^{***}	
	(0.28)	(0.29)	(0.34)	(0.31)	
Nr of Obs.	$147,\!420$	$134,\!295$	$92,\!826$	$131,\!164$	

Panel A: Unemployed Routine Manual \rightarrow Employed Routine Manual ($URM \rightarrow ERM$)

Panel B: Not in the Labor Force \rightarrow Employed Routine Manual	$(NLF \rightarrow$	ERM)
--	--------------------	------

	Baseline Expansion $(1976m1-1979m12)$: 1.23%					
-	1982m12-	1991m4-	2001m12-	2009m7-		
	1990m6	2001m2	2007m11	2018m12		
Total Change	-0.12^{***}	-0.09^{***}	-0.13^{***}	-0.35^{***}		
	(0.02)	(0.02)	(0.02)	(0.02)		
Composition	-0.02^{***}	$+0.05^{***}$	$+0.11^{***}$	$+0.11^{***}$		
	(0.00)	(0.01)	(0.01)	(0.01)		
Propensities	-0.11^{***}	-0.14^{***}	-0.24^{***}	-0.46^{***}		
	(0.02)	(0.02)	(0.02)	(0.02)		
Nr of Obs.	$2,\!933,\!617$	$3,\!127,\!397$	2,347,264	$3,\!375,\!160$		

Note: The numbers represent percentage point changes. The Composition component corresponds to the change explained by demographic characteristics (age, education, gender, race), while the Propensities component is driven by changes in estimated coefficients (changes in estimated transition probabilities, conditional on demographic characteristics). Standard errors are adjusted to account for clustering at the individual level. * : p < 0.10, ** : p < 0.05, *** : p < 0.01

Panel B shows that the inflow rate from non-participation to routine cognitive employment actually increases during the 1990s and 2000s expansions, relative to the pre-polarization rate of the 1980s. There is, however, a precipitous fall in this transition rate after the Great Recession. Note that during all of the polarization era expansions, changes in propensities contribute to a decline in the transition rate from non-participation to routine cognitive employment, indicating that from the 1990s onwards, the average non-participant becomes less likely to transition to routine cognitive employment, conditional on demographic characteristics.

Crucially, Appendix Table A.2 shows that similar declines in inflow rates are *not* observed when considering transitions into *non*-routine occupations. Hence, the decline in the inflow rates to routine occupations is *not* driven by an economy-wide decline in inflows to employment, but rather, by changes that are specific to routine occupations.

	Baseline Expansion ($1982m12-1990m6$): 15.77%				
	1991m4-	2001m12-	2009m7-		
	2001m2	2007m11	2018m12		
Total Change	-0.31	-2.36^{***}	-6.25^{***}		
	(0.24)	(0.26)	(0.22)		
Composition	-0.19^{***}	-0.45^{***}	-0.54^{***}		
	(0.05)	(0.07)	(0.07)		
Propensities	-0.12	-1.91^{***}	-5.71^{***}		
	(0.24)	(0.27)	(0.23)		
Nr of Obs.	120,754	97,892	$128,\!564$		

Panel A: Unemployed Routine Cognitive \rightarrow Employed Routine Cognitive ($URC \rightarrow ERC$)

Panel Ba	: Not in	the Labo	r Force \rightarrow	· Employed	Routine	Cognitive	(NLF)	$\rightarrow ERC$)
						()	\ \	

	Baseline Expansion $(1982m12-1990m6): 1.58\%$				
	1991m4-	2001m12-	2009m7-		
	2001 m 2	2007 m 11	2018m12		
Total Change	+0.02	$+0.05^{***}$	-0.34^{***}		
	(0.02)	(0.02)	(0.01)		
Composition	$+0.08^{***}$	$+0.21^{***}$	$+0.20^{***}$		
	(0.00)	(0.01)	(0.01)		
Propensities	-0.06^{***}	-0.16^{***}	-0.54^{***}		
	(0.01)	(0.02)	(0.01)		
Nr of Obs.	4,281,880	$3,\!501,\!747$	$4,\!529,\!643$		

Note: The numbers represent percentage point changes. The Composition component corresponds to the change explained by demographic characteristics (age, education, gender, race), while the Propensities component is driven by changes in estimated coefficients (changes in estimated transition probabilities, conditional on demographic characteristics). Standard errors are adjusted to account for clustering at the individual level. * : p < 0.10, *** : p < 0.05, *** : p < 0.01

4.2 Demographics-Based Counterfactual

We return to our counterfactual exercises in order to assess the role of demographic change and of changes in transition propensities. Section 3 presented results from counterfactuals holding *aggregate* transition rates constant. This is akin to holding *both* the demographic composition and the group-specific transition propensities constant. Specifically, consider the law-of-motion from Equation (2). The vector of stocks at any time t can be written as a sum over J demographic group-specific stocks:

$$Stocks_t = \sum_{j=1}^{J} \Omega_{jt} Stocks_{jt}$$
⁽⁷⁾

where j denotes demographic groups, Ω_{jt} is the share of the total population in the economy at time t belonging to group j, and $Stocks_{jt}$ is a vector with the share of individuals from demographic group j within each of the nine labor market states. Naturally $\Sigma_{j=1}^{J}\Omega_{j,t} = 1$. Note also that the evolution of the vector of demographic group-specific stocks, $Stocks_{jt}$, can be described through a law of motion analogous to Equation (2):

$$Stocks_{j,t+1} = \rho_{jt} * Stocks_{jt} \tag{8}$$

where ρ_{jt} is a (9 x 9) transition rate matrix analogous to the one in Equation (1), but with transition rates specific to group j. It follows that the aggregate law of motion can be re-written as:

$$Stocks_{t+1} = \sum_{j=1}^{J} \Omega_{j,t+1} * \rho_{jt} * Stocks_{jt}$$

$$\tag{9}$$

Even if transition propensities had remained constant for all demographic groups, the *aggregate* transition rates and stocks would have evolved over time due to changes in the demographic composition of the U.S. economy. By holding certain aggregate transition rates constant at pre-polarization levels in our counterfactual exercises in Section 3, we effectively abstracted from the impact of demographic changes such as the increased educational attainment and aging experienced in the U.S. population.

On the other hand, while the OB decomposition analysis of Section 4.1 helps us understand the role of demographic change relative to propensity changes, it ignores the fact that, had transition propensities remained at their phase 0 levels, the demographic composition of labor market state A in phase 1 would potentially differ dramatically from that in the data. Specifically, the term in the OB decomposition attributed to demographics is given by:

$$\left(\overline{X}_0^A - \overline{X}_1^A\right)\widehat{\beta}_0$$

This is the difference between the observed aggregate transition rate in period 0, $(\overline{X}_0^A \hat{\beta}_0)$, and the counterfactual rate that would be observed using the transition propensities from period 0 and the *observed* demographic composition of period 1, $(\overline{X}_1^A \beta_0)$. It abstracts from the fact that the counterfactual demographic composition of period 1 would be different if the transition propensities had remained at their period 0 levels.

In this section, we perform counterfactuals that account for the fact that the composition in different labor market states evolves endogenously according to changes in different groups' transition rates. This provides a more accurate assessment of the importance of demographic and propensity changes.

In each period, we divide individuals into 36 demographic groups according to their gender, age, and education (2 gender groups \times 3 education groups \times 6 age groups). We then calculate the time series of transition rates across the nine labor market states for each demographic group, i.e. the matrix ρ_{jt} from Equations (8) and (9). For each of the 36 groups, we track their distribution across labor market states over time using either true or counterfactual transition rates by applying the law-of-motion from Equation (8). This provides a labor market evolution for each demographic group that is consistent with the transition rates being considered.

We then apply the aggregation in Equation (7) in order to obtain the aggregate stock in each labor market state at each point in time. Each demographic group's weight Ω_{jt} is equal to the group's *share* of the total population, as observed in the data. This weighting procedure ensures that we match the evolution of the aggregate demographic composition over time, while simultaneously ensuring that the distribution of each group *across* the nine labor market states is determined endogenously.²² An analogous interpretation of our weighting approach is that entry and exit from the sample occurs in proportion to the size of each labor market state within a demographic bin, so that it does not change *group-specific* labor force composition. However, different entry and exit rates across groups change their relative population size, thus changing the labor market composition *in the aggregate*.

4.2.1 The role of demographic change

We begin by quantifying the overall role of demographic composition change in the U.S. population for the decline in routine employment. We do this by holding the demographic composition constant at pre-polarization levels while allowing the transition rates of each demographic group to evolve over time as observed in the data. Specifically, we determine counterfactual stocks by modifying Equation (9) and keeping the weight of the different demographic groups constant. Thus, the demographic counterfactual becomes:

$$Stocks_{t+1}^{CF} = \sum_{j=1}^{J} \Omega_j^{CF} * \rho_{jt} * Stocks_{jt},$$
(10)

where Ω_j^{CF} is group j's share of the total population in the first observation in our data, January 1976. By holding the demographic structure constant, this removes any changes that are driven by the changing relative size of groups in the economy (due to rising educational attainment, or population aging, for example), while still allowing changes within groups to occur as in the data. Hence, any decline in routine employment mitigated by this counterfactual is solely due to demographic change. As before, we summarize the results from the counterfactual experiment by computing the fraction of the decline in routine employment (from its peak to its value at the end of the sample, period T), that is avoided in the counterfactual, e.g.:

$$1 - \frac{\Delta ERM^{CF}}{\Delta ERM} = 1 - \frac{ERM_T^{CF} - ERM_{Peak}}{ERM_T - ERM_{Peak}}$$

The first line of Table 6 presents the results for this counterfactual exercise. We find that, relative to the peak of ERM, holding demographics constant mitigates slightly less than a quarter of its fall. By contrast, this demographic counterfactual mitigates none of the observed fall for ERC.

4.2.2 The role of inflows and outflows to routine employment

Our next set of exercises holds the transition rates of interest constant at their pre-job polarization levels for all 36 demographic groups; these are the inflow rates to routine employment from non-participation and unemployment, and the outflow rates from routine employment to non-participation and unemployment. All other group-specific transition rates and the demographic composition of the population are allowed to evolve as in the data. Thus, vis-a-vis the analysis in Section 3, the key transition rates will evolve over time *in the aggregate* due to changes in demographic composition; we only remove the effects due to *group-specific* transition propensity changes.

 $^{^{22}}$ One might worry that using a large number of demographic bins leads to noisy group-specific transition rates. We verify that employment in routine occupations constructed using group-specific transition rates along with the weighting procedure tracks the series constructed from aggregate transition rates extremely closely. We use the series derived from the demographic group-specific rates as our benchmark throughout this section of the paper.

	Fraction of decline/increase avoided				
	Employed: Routine	Employed: Routine	Not in the		
	Manual (ERM)	Cognitive (ERC)	Labor Force (NLF)		
Demographics	0.23	-0.03			
Inflow and Outflow Rates	0.40	0.44	0.75		
Inflow Rates	0.26	0.45	0.60		
Inflow from Non-Participation	0.16	0.30	0.46		
Inflow from Unemployment	0.11	0.16	0.16		
Outflow rates	0.16	0.01	0.16		
Outflow to Non-Participation	0.11	0.16	0.09		
Outflow to Unemployment	0.05	0.03	0.06		

Table 6: Results from counterfactual exercises, allowing for demographic group heterogeneity

Note: The table reports the fraction of the decline in ERM and ERC (from their respective peaks to their levels in 2018:12), and the fraction of the rise in NLF (from its level in the early 2000s to its level in 2018:12), that is avoided by holding either demographics, or different sets of transition rates constant at their pre-polarization levels within demographic groups. The ERM, ERC and NLF stocks are measured as a fraction of the working-age population.

More formally, we compute counterfactual stocks based on:

$$Stocks_{t+1}^{CF} = \sum_{j=1}^{J} \Omega_{j,t+1} * \rho_{jt}^{CF} * Stocks_{jt}^{CF}$$
(11)

where ρ_{jt}^{CF} is a counterfactual matrix of (group-specific) transition rates, and, as before, in the initial period we use the observed stocks in the data: $Stocks_{j0}^{CF} = Stocks_{j0}$. The procedure used to compute ρ_{jt}^{CF} is analogous to Section 3, the only difference being that the counterfactual transition rate matrix is now computed for each demographic group j. We allow for changes in demographic composition by weighting each group based on its observed share of the total population in each period, Ω_{jt} .

The results for the fraction of the decline in routine employment avoided in each counterfactual, again computed as in Equation (4), are presented in the remainder of Table 6. The results are quite similar to the ones obtained based on the analysis of aggregate transition rates in Section 3. Jointly, the inflow and outflow transition rates account for about 40 percent of the fall in *ERC* and *ERM*. The inflow rates play a more important role than the outflow rates and, again, the inflow rates from non-participation into routine employment are more important than the inflow rates from unemployment. Interestingly, we find a slightly more important role for outflow rates in accounting for the decline in *ERM*, compared to the results obtained in Table 3.

4.2.3 Implications for non-participation

The ability to decompose demographic changes from group-specific changes allows us to address other important developments in the U.S. labor market. In particular, how much of the increase in labor force non-participation can be accounted for by group-specific changes in these same transition rates to and from routine employment? This is considered in the rightmost column of Table 6. We focus on how much of the increase in non-participation between the early 2000s (where it accounted for 31.2% of the

population) to its level at the end of 2018 (35.9%) is mitigated in the different counterfactual exercises.²³

Inflow and outflow rates to and from routine employment jointly account for about three quarters of the overall rise in labor force non-participation. As with the fall in routine employment, the inflow rates account for the bulk of the change. Specifically, changes in inflow rates to routine employment account for 60 percent of the overall rise in non-participation; changes in the outflow rates from these occupations account for 16 percent.

Overall, a consistent result emerges: the fall in inflow rates, and mainly those from non-participation, accounts for the bulk of the fall in routine employment. This same change accounts for a substantial proportion of the increase in non-participation observed since the turn of the century.

4.3 Which are the "key" demographic groups?

So far, we have highlighted the overall importance of changes in transition propensities in the decline of routine employment. The changes in transition propensities over the past 40 years would have been experienced differentially across demographic groups. In this section, we investigate which specific groups' transition rate changes account for the bulk of the *aggregate* changes that we have documented.

To do this, we recompute our previous counterfactuals, but this time holding rates constant only for specific demographic groups. For instance, to isolate the role of propensity changes of males, we hold constant the transition rates only for the 18 (out of 36) demographic bins belonging to men, allowing the rates of women to evolve as in the data. We partition the demographic groups into J^A and J^B , where $J^A + J^B = J$, and use the observed transition rate matrix ρ_{jt} for groups in J^A , and the counterfactual transition rate matrix ρ_{jt}^{CF} for groups in J^B . The resulting counterfactual stocks evolve according to:

$$Stocks_{t+1}^{CF} = \sum_{j \in J^A} \Omega_{j,t+1} * \rho_{jt} * Stocks_{jt} + \sum_{j \in J^B} \Omega_{j,t+1} * \rho_{jt}^{CF} * Stocks_{jt}^{CF}$$
(12)

We perform this exercise along the three dimensions that characterize our groups: (i) gender (male, female), (ii) education (less than high school, high school diploma or some college, college graduates), and (iii) age (16-34 year olds, 35-54 year olds, over 55). Table 7 reports the results. The numbers in the table report the fraction of the *overall* change that can be avoided by "freezing" the transition rates for the specific demographic group in question; these fractions are once again computed as in Equation (4).

We note several results. First, the quantitative role of the inflow rate is more important than the outflow rates. The only demographic split where outflow rates matter more are for middle-age transitions to and from ERM. This reconfirms our earlier results.

Second, by holding rates constant for males only, we account for roughly a third of the overall fall in routine manual employment, and more than half of the rise in non-participation. On the other hand, it is changes in the inflow rates to routine cognitive among women that account for the bulk of the employment decline in these occupations. Third, along the age dimension, changes among the young (16-34 year olds) account for almost 40% of the fall in ERM and slightly more than half of the rise in overall non-participation.²⁴ Fourth, changes in transitions propensities for those with intermediate

 $^{^{23}}$ This is computed in an analogous way to routine employment, using Equation (4), but with NLF as our stock of interest.

 $^{^{24}}$ These findings are in line with the evidence presented in Beaudry et al. (2014) and Beaudry et al. (2016), who show that there are important changes in the occupational composition of employment for young workers since the 1990s.

Fraction of aggregate decline/increase avoided										
Panel A: Gender		Males]	Females				
	ERM	ERC	NLF		ERM	ERC	NLF			
Inflow and Outflow Rates	0.38	0.15	0.54		0.02	0.29	0.21			
Inflow Rates	0.23	0.09	0.32		0.03	0.36	0.29			
Outflow Rates	0.16	0.06	0.24		-0.01	-0.07	-0.08			
Panel B: Age		16 - 34				35 - 54			55 - 75	
	ERM	ERC	NLF		ERM	ERC	NLF	ERM	ERC	NLF
Inflow and Outflow Rates	0.39	0.27	0.52		0.12	0.16	0.23	0.01	0.00	-0.00
Inflow Rates	0.21	0.17	0.36		0.04	0.18	0.16	0.01	0.09	0.09
Outflow Rates	0.06	0.10	0.16		0.09	-0.01	0.09	0.00	-0.10	-0.08
Panel C: Education	Les	s than 1	HS		HS -	+ Some	Col	Co	llege Gr	ad
	ERM	ERC	NLF		ERM	ERC	NLF	ERM	ERC	NLF
Inflow and Outflow Rates	0.11	0.04	0.22		0.29	0.31	0.52	0.00	0.08	0.03
Inflow Rates	0.06	0.03	0.13		0.18	0.26	0.38	0.01	0.15	0.09
Outflow Rates	0.04	0.01	0.07		0.12	0.05	0.16	-0.01	-0.06	-0.07

Table 7: Results from counterfactual exercises focusing on specific demographic groups

Note: The table reports the fraction of the aggregate decline in ERM and ERC (from their respective peaks to their levels in 2018:12), and the fraction of the rise in NLF, that is avoided by holding different sets of transition rates constant at their pre-polarization levels for specific demographic groups. The ERM, ERC and NLF stocks are measured as a fraction of the working-age population.

levels of education drive the majority of the changes in routine employment and non-participation in our counterfactuals. Clearly, for these demographic groups (men, the young, and the intermediately educated), falling inflow rates to routine employment have not been matched by increasing inflow rates to non-routine occupations; rather, they have resulted in increased propensities to remain non-employed.

Appendix Figures A.1 through A.4 depict the average inflow rates to routine employment across the six expansionary phases for each of the demographic groups that we consider. These confirm the strong declines in the inflow rates to ERM experienced by men, particularly young men, as well as the decline in the inflow rates to ERC among women, particularly those with intermediate and higher levels of education.

5 Conclusions

We analyze changes in worker flows that account for the decline in routine employment in the U.S. economy, using matched individual-level data from the monthly CPS. Quantitatively, decreases in inflow rates to routine employment (from unemployment and non-participation) play a much larger role than increases in outflow rates. Changes in aggregate transition rates are primarily driven by changes within demographic groups, rather than changes in the demographic composition of the U.S. population. We find that changes in the transition propensities of males, young individuals, and those with intermediate levels of education are of primary importance. Moreover, changes in inflow and outflow rates to/from routine employment among these groups also account for a substantial portion of the rise in labor force non-participation observed in recent decades. Declining inflow rates to routine occupations for these groups have not been accompanied by increasing inflow rates to non-routine occupations, in spite of

non-routine occupational growth in the aggregate.

Our findings provide a richer picture of the way polarization has occurred over recent decades, and provide guidance for policy and the equilibrium models needed to inform its formulation. For instance, our results show that polarization and declining labor force participation are related phenomena, and highlight the importance of changing transition rates between non-participation and routine employment. There exists a well-developed literature on frictional labor market dynamics, studying employmentunemployment flows among labor force participants (see Diamond (1982), Mortensen (1982), Pissarides (1985), and the subsequent literature). Much less has been done to model the flows between participation and non-participation, and more work along such lines is warranted (see, for instance, Krusell et al. (2011)).

With regard to routine employment, a concern among policymakers is that those employed in such middle-wage occupations are largely prime-aged or older, and face "job displacement" risk (see, for instance, Jacobson et al. (1993) and the subsequent literature) due to elevated separation rates. Although certainly valid, we find that (aside from normal cyclical spikes in employment separations at the onset of recessions) outflow rates from routine employment during economic expansions are largely unchanged. Instead, declining routine employment overall is largely due to declining inflow rates for those who used to find employment in these occupations. Moreover, these deteriorating labor market prospects are more acute for younger workers as opposed to older ones. Improving employment prospects of such individuals through retraining programs and interventions that do not necessarily prescribe the attainment of a college degree is of high priority (see, for instance, Holzer (2015) and Jaimovich et al. (2020)). Further analysis of the efficacy of specific active labor market interventions along the lines of Card et al. (2018) is always warranted. Finding innovative solutions for these worker groups, given the diminished opportunities in routine occupations, is of first order importance.

References

- Acemoglu, Daron and David Autor (2011), "Skills, tasks and technologies: Implications for employment and earnings." *Handbook of Labor Economics*, 4, 1043–1171.
- Autor, David and David Dorn (2009), "This job is "getting old": Measuring changes in job opportunities using occupational age structure." American Economic Review, 99, 45–51.
- Autor, David H. and David Dorn (2013), "The growth of low skill service jobs and the polarization of the U.S. labor market." *American Economic Review*, 103, 1553–1597.
- Autor, David H, David Dorn, and Gordon H Hanson (2015), "Untangling trade and technology: Evidence from local labor markets." *Economic Journal*, 125, 621–646.
- Autor, D.H., L.F. Katz, and M.S. Kearney (2006), "The Polarization of the US Labor Market." The American Economic Review, 96, 189–194.
- 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, Paul, David A. Green, and Benjamin M. Sand (2014), "The declining fortunes of the young since 2000." American Economic Review, 104, 381–86.
- Beaudry, Paul, David A Green, and Benjamin M Sand (2016), "The great reversal in the demand for skill and cognitive tasks." *Journal of Labor Economics*, 34, S199–S247.
- Blinder, Alan (1973), "Wage discrimination: Reduced form and structural estimates." Journal of Human Resources, 8, 436–455.
- Card, David, Jochen Kluve, and Andrea Weber (2018), "What works? A meta analysis of recent active labor market program evaluations." *Journal of the European Economic Association*, 16, 894–931.
- Christiano, Lawrence J and Terry J Fitzgerald (2003), "The band pass filter." International Economic Review, 44, 435–465.
- Cortes, G.M. (2016), "Where have the middle-wage workers gone? A study of polarization using panel data." Journal of Labor Economics, 34, 63–105.
- Cortes, Guido Matias, Nir Jaimovich, and Henry E. Siu (2017), "Disappearing routine jobs: Who, how and why?" *Journal of Monetary Economics*, 91, 69–87.
- Darby, Michael R., John Haltiwanger, and Mark Plant (1985), "Unemployment rate dynamics and persistent unemployment under rational expectations." The American Economic Review, 75, 614–637.
- Diamond, Peter A (1982), "Wage determination and efficiency in search equilibrium." The Review of Economic Studies, 49, 217–227.
- Elsby, M.W.L., R. Michaels, and G. Solon (2009), "The ins and outs of cyclical unemployment." American Economic Journal: Macroeconomics, 1, 84–110.

- Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux (2011), "Occupational tasks and changes in the wage structure." University of British Columbia.
- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren (2018), "Integrated Public Use Microdata Series (IPUMS), Current Population Survey: Version 6.0 [dataset]."
- Foote, Christopher L. and Richard W. Ryan (2014), "Labor-market polarization over the business cycle." In NBER Macroeconomics Annual 2014, Volume 29, 371–413, University of Chicago Press.
- Fortin, Nicole, Thomas Lemieux, and Sergio Firpo (2011), "Decomposition methods in economics." Handbook of Labor Economics, 4, 1–102.
- Frazis, Harley J, Edwin L Robison, Thomas D Evans, and Martha A Duff (2005), "Estimating gross flows consistent with stocks in the CPS." *Monthly Labor Review*, 128, 3.
- Goos, M. and A. Manning (2007), "Lousy and lovely jobs: The rising polarization of work in Britain." The Review of Economics and Statistics, 89, 118–133.
- Goos, M., A. Manning, and A. Salomons (2014), "Explaining job polarization: Routine-biased technological change and offshoring." *American Economic Review*, 104, 2509–2526.
- Goos, Maarten, Alan Manning, and Anna Salomons (2009), "Job polarization in Europe." American Economic Review, 99, 58–63.
- Hall, Robert E (2006), "Job loss, job finding and unemployment in the US economy over the past 50 years." In NBER Macroeconomics Annual 2005, Volume 20, 101–166, MIT Press.
- Holzer, Harry (2015), "Job market polarization and US worker skills: A tale of two middles." Brookings Institution, Economic Studies, policy brief.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan (1993), "Earnings losses of displaced workers." American Economic Review, 83, 685–709.
- Jaimovich, N. and H.E. Siu (2012), "The trend is the cycle: Job polarization and jobless recoveries." National Bureau of Economic Research Working Paper.
- Jaimovich, Nir, Itay Saporta-Eksten, Henry E. Siu, and Yaniv Yedid-Levi (2020), "The macroeconomics of automation: Data, theory, and policy analysis." University of Zurich working paper.
- Kambourov, G. and I. Manovskii (2013), "A Cautionary Note on Using (March) CPS and PSID Data to Study Worker Mobility." *Macroeconomic Dynamics*, 17, 172–194.
- Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Aysegul Sahin (2011), "A three state model of worker flows in general equilibrium." *Journal of Economic Theory*, 146, 1107–1133.
- Meyer, PB and AM Osborne (2005), "Proposed Category System for 1960-2000 Census Occupations." Bureau of Labor Statistics Working Paper, 383.
- Mortensen, Dale T (1982), "The matching process as a noncooperative bargaining game." In *The eco*nomics of information and uncertainty, 233–258, University of Chicago Press.

- Moscarini, Giuseppe and Kaj Thomsson (2007), "Occupational and job mobility in the US." Scandinavian Journal of Economics, 109, 807–836.
- Oaxaca, Ronald (1973), "Male-female wage differentials in urban labor markets." *International Economic Review*, 14, 693–709.
- Pissarides, Christopher A. (1985), "Short-run equilibrium dynamics of unemployment, vacancies, and real wages." *The American Economic Review*, 75, 676–690.
- Shimer, Robert (2012), "Reassessing the ins and outs of unemployment." *Review of Economic Dynamics*, 15, 127–148.
- Smith, Christopher L. (2013), "The dynamics of labor market polarization." Finance and Economics Discussion Series 2013-57. Board of Governors of the Federal Reserve System.

Appendix A Additional Tables and Figures

Broad	Census Coding System				
Occupation	1970	1980 and 1990	2002	2010	
Non-Routine	001-100, 102-162, 165,	003-225, 228-	0010-3540	0010-3540	
Cognitive	171, 173-216, 222-225,	229, 234-235,			
	230, 235-245, 321, 326,	473-476			
	363, 382, 426, 506, 801-				
	802, 924, 926				
Non-Routine	101, 505, 740, 755, 901-	403-469, 486-	3600 - 4650	3600 - 4650	
Manual	923, 925, 931-984	487,773			
		2 12 222			
Routine	220, 231-233, 260-285, 201, 205, 200, 200, 200, 200, 200, 200, 200	243-389	4700-5930	4700-5940	
Cognitive	301-305, 310-320, 323-				
	325, 330-362, 364-381,				
	383-395				
Boutine	163-164 170 179 221	<u>996_997_933</u>	6200-9750	6200-9750	
Manual	226 401-425 430-446	503_760 774_	0200-5150	0200-3150	
manuai	45250, 401-425, 450-440, 452504, 510, 575, 601	700 803 860			
	624 626 715 750 751	873 880			
	753 754 760 762 785	010-009			
	105-104, 100, 102-105				
Farming.	450, 580, 600, 625, 752,	477-485, 488-	6000-6130.	6005-6130.	
Military	761, 821-824	499, 905	9800-9840	9800-9840	

Table A.1: Mapping of detailed occupation codes to broad groups

Table A.2: Oaxaca Decomposition: Inflows to Non-Routine Occupations

Panel A: $UNRC \rightarrow ENRC$

i unoi 11. 0 10100	/ 11/100				
	Baseline Expansion $(1982m12-1990m6): 14.72\%$				
	1991 m 4 - 2001 m 2	2001m12-2007m11	2009m7-2018m12		
Total Change	$+1.45^{***}$	$+1.40^{***}$	-0.45		
	(0.22)	(0.36)	(0.30)		
Composition	+0.12	$+0.49^{***}$	$+0.63^{***}$		
	(0.09)	(0.13)	(0.11)		
Propensities	$+1.33^{***}$	$+0.91^{**}$	-1.08^{***}		
	(0.33)	(0.36)	(0.31)		
Nr of Obs.	66,747	$58,\!272$	$91,\!379$		

Panel B: $NLF \rightarrow ENRC$

	Baseline Expansion (1982m12-1990m6): 0.97%				
	1991m4-2001m2	2001m12-2007m11	2009m7-2018m12		
Total Change	$+0.16^{***}$	$+0.38^{***}$	$+0.33^{***}$		
	(0.01)	(0.01)	(0.01)		
Composition	$+0.16^{***}$	$+0.37^{***}$	$+0.44^{***}$		
	(0.00)	(0.01)	(0.01)		
Propensities	0.00	0.00	-0.11^{***}		
	(0.01)	(0.01)	(0.01)		
Nr of Obs.	4,281,880	$3,\!501,\!747$	4,529,643		

Panel C: $UNRM \rightarrow ENRM$

	Baseline Expansion $(1976m1-1979m12): 14.15\%$				
	1982m12-1990m6	1991m4-2001m2	2001m12-2007m11	2009m7-2018m12	
Total Change	-0.75^{**}	$+1.41^{***}$	$+1.33^{***}$	-1.22^{***}	
	(0.32)	(0.33)	(0.35)	(0.31)	
Composition	-0.41^{***}	-0.20^{**}	-0.06	-0.54^{***}	
	(0.07)	(0.09)	(0.12)	(0.13)	
Propensities	-0.34	$+1.60^{***}$	$+1.39^{***}$	-0.68^{**}	
	(0.33)	(0.34)	(0.37)	(0.34)	
Nr of Obs.	70,118	$72,\!296$	$54,\!627$	87,066	

Panel D: $NLF \rightarrow ENRM$

	Baseline Expansion $(1976m1-1979m12): 1.59\%$				
	1982m12-1990m6	1991m4-2001m2	2001m12-2007m11	2009m7-2018m12	
Total Change	-0.09^{***}	-0.11^{***}	-0.03^{*}	-0.23^{***}	
	(0.02)	(0.02)	(0.02)	(0.02)	
Composition	-0.15^{***}	-0.16^{***}	-0.10^{***}	-0.09^{***}	
	(0.00)	(0.01)	(0.01)	(0.01)	
Propensities	$+0.06^{***}$	$+0.04^{**}$	$+0.07^{***}$	-0.14^{***}	
	(0.02)	(0.02)	(0.02)	(0.02)	
Nr of Obs.	$2,\!933,\!617$	$3,\!127,\!397$	2,347,264	$3,\!375,\!160$	

Note: The numbers represent percentage point changes. The Composition component corresponds to the change explained by demographic characteristics (age, education, gender, race), while the Propensities component is driven by changes in estimated coefficients (changes in estimated transition probabilities, conditional on demographic characteristics). Standard errors are adjusted to account for clustering at the individual level. * : p < 0.10, ** : p < 0.05, *** : p < 0.01



Figure A.1: Group-Specific Transition Rates: $NLF \rightarrow ERM$

Note: Each panel displays the average transition rate for a specific demographic group across our six expansionary phases, E1–E6 (see Table 2 for details). Each group is identified by a three character index: the first character denotes the gender (M for Males and F for Females), the second one denotes the age groups (in increasing order from 1 to 6), and the last one identifies the educational attainment (L=high school dropout, M=high school graduate and some college, H=college graduates and above).



Figure A.2: Group-Specific Transition Rates: $URM \rightarrow ERM$

Note: Each panel displays the average transition rate for a specific demographic group across our six expansionary phases, E1–E6 (see Table 2 for details). Each group is identified by a three character index: the first character denotes the gender (M for Males and F for Females), the second one denotes the age groups (in increasing order from 1 to 6), and the last one identifies the educational attainment (L=high school dropout, M=high school graduate and some college, H=college graduates and above).



Figure A.3: Group-Specific Transition Rates: $NLF \rightarrow ERC$

Note: Each panel displays the average transition rate for a specific demographic group across our six expansionary phases, E1–E6 (see Table 2 for details). Each group is identified by a three character index: the first character denotes the gender (M for Males and F for Females), the second one denotes the age groups (in increasing order from 1 to 6), and the last one identifies the educational attainment (L=high school dropout, M=high school graduate and some college, H=college graduates and above).



Figure A.4: Group-Specific Transition Rates: $URC \rightarrow ERC$

Note: Each panel displays the average transition rate for a specific demographic group across our six expansionary phases, E1–E6 (see Table 2 for details). Each group is identified by a three character index: the first character denotes the gender (M for Males and F for Females), the second one denotes the age groups (in increasing order from 1 to 6), and the last one identifies the educational attainment (L=high school dropout, M=high school graduate and some college, H=college graduates and above).