On the Demographic Adjustment of Unemployment *

Regis Barnichon† and Geert Mesters‡

July 17, 2016

Abstract

The unemployment rate is one of the most important business cycle indicators, but its interpretation can be difficult because slow changes in the demographic composition of the labor force affect the level of unemployment and make comparisons across business cycles difficult. To purge the unemployment rate from demographic factors, labor force shares are routinely used to control for compositional changes. This paper shows that this approach is ill-defined, because the labor force share of a demographic group is mechanically linked to that group's unemployment rate, as both variables are driven by the same underlying worker flows. We propose a new demographic-adjustment procedure that uses a dynamic factor model for the worker flows to separate aggregate labor market forces and demographic-specific trends. Using the US labor market as an illustration, our demographic-adjusted unemployment rate indicates that the 2008-2009 recession was much more severe and generated substantially more slack than the early 80s recession.

JEL classification: J6, E24

Keywords: accounting, demographics, stock-flow, unobserved components, dynamic factor models.

*We would like to think Tomaz Cajner, Bart Hobijn, Kris Nimark, Yanos Zylberberg and seminar participants for helpful comments and Roger Gomis for excellent research assistance. † Regis Barnichon: CREI, Universitat Pompeu Fabra and CEPR, email: rbarnichon@crei.cat. Barnichon acknowledges financial support from the Spanish Ministerio de Economia y Competitividad (grant ECO2011-23188), the Generalitat de Catalunya (grant 2009SGR1157) and the Barcelona GSE Research Network. ‡ Geert Mesters: Universitat Pompeu Fabra, Barcelona GSE and The Netherlands Institute for the Study of Crime and Law Enforcement, email: geert.mesters@upf.edu. Mesters acknowledges support from the Marie Curie FP7-PEOPLE-2012-COFUND Action. Grant agreement no: 600387. Any errors are our own.

1 Introduction

The unemployment rate is one of the most important business cycle indicators. For instance, Hatzius & Stehn (2012) of Goldman Sachs refer to the unemployment rate as their "desert island economic indicator", the one they would choose if they had to choose only one indicator to provide information about the economy.

However, despite its appeal as a cyclical indicator, the specific level of the unemployment rate can be difficult to interpret –and the state of the business cycle difficult to assess–, because slow changes in the demographic composition of the labor force affect the level of unemployment and make comparisons across business cycles difficult.[1]

To address this issue and purge the unemployment rate from such compositional effects, researchers, policy makers and practitioners typically rely on a shift-share analysis where labor force shares are used to "control" for composition.[2] In particular, a demographic-adjusted unemployment rate is often constructed by holding the labor force shares fixed and letting only the group-specific unemployment rates fluctuate.

In this paper, we show that such a "labor force shift-share" (LFSS) analysis does not appropriately adjust the aggregate unemployment rate for demographic changes. A LFSS analysis is biased, because it does not take into account the fact that a group's labor force share is *mechanically* linked to that group's unemployment rate, as both variables are driven by the same underlying worker flows. To see that, recall that the unemployment rate and the participation rate of each demographic group are stocks determined by the flows of workers in and out of employment, unemployment and nonparticipation.[3] But since the labor force share of any demographic group depends on that group's participation rate, movements in the labor force share are not independent from movements in the group's unemployment rate. As a result, one cannot, as in a LFSS approach, hold the labor force share fixed while

[1]For instance, because younger workers have a higher unemployment rate than older workers, an older population will mechanically have a lower aggregate unemployment rate.

[2]See for instance Perry (1970), Flaim (1979), Gordon (1998), Summers (1986), Shimer (1998), Horn & Heap (1999), Katz & Krueger (1999), Shimer (2001), Abraham & Shimer (2001), Herbertsson, Gylfi & Phelps (2002), Valletta & Hodges (2005), Kudlyak, Reilly & Slivinski (2010), Elsby, Hobijn & Sahin (2010), Lazear & Spletzer (2012) as well as numerous policy reports (e.g., the CEA Economic Report to the President, 1997).

[3]A large literature (going back to Kaitz (1970), Perry (1972), Blanchard & Diamond (1990) and more recently Shimer (2012) and Elsby, Hobijn & Sahin (2013)) has shown that the labor market is characterized by large worker flows taking place between employment, unemployment and nonparticipation, and that the unemployment and participation rates are stocks determined by these underlying worker flows.

letting the unemployment rate fluctuate (or vice-versa).

We propose a new demographic-adjustment procedure. The idea is to construct a counter-factual aggregate unemployment rate that is driven solely by aggregate labor market forces and not by demographic group-specific changes. To do so, we use a dynamic factor model to identify common variations across demographic groups and extract the common factors affecting the labor market.[4] Importantly, we work directly with the worker flows instead of working with the unemployment and participation stocks. With three labor market states –Employment, Unemployment and Nonparticipation–, there are six possible transitions, and we use a dynamic factor model for each panel of worker flows to decompose the worker flows of the different demographic groups into a common component and several demographic-specific components.[5] The common components are then used in a stock-flow model of the labor market to construct a demographic-adjusted unemployment rate.

We illustrate our approach with the demographic adjustment of the US unemployment rate. Our demographic-adjusted unemployment rate indicates that the three labor market peaks of 1979, 1989 and 2006 were remarkably similar. In all cases the (counter-factual) unemployment rate came down to about 5 percent. The 2001 labor market peak appears different however (with an unemployment rate as low as 4.2 percent), in line with the common perception that the high-tech boom led to an exceptional expansionary phase. Comparing troughs, our demographic-adjusted unemployment rate indicates that the 2008-2009 recession was much more severe and generated substantially more slack than the early 80s recession. This contrasts with a LFSS-adjusted unemployment rate, which points to similar amounts of slack across recessions.

Overall, our method indicates that demographic changes lowered the aggregate unemployment rate by about 2.2 percentage points (ppt) since the mid 70s. This contribution is the sum of two separate effects: (i) changes in the population structure, due to popula-

[4]Our use of a factor model to isolate common labor market forces echoes (albeit in a very different context) recent approaches aimed at measuring global business cycles (see Kose, Otrok & Whiteman (2003), Kose, Otrok & Whiteman (2008) and Del Negro & Otrok (2008)) and trend inflation (see Stock & Watson (2015)). We adopt a small scale factor model in the spirit of Engle & Watson (1981, 1983), Stock & Watson (1989), Quah & Sargent (1993) and Otrok & Whiteman (1998). Such factor models are also often referred to as multivariate unobserved components models.

[5]Specifically, for each A-B transition (i.e., from state A to state B), we have a panel of A-B flows made of the I demographic groups, and we use a dynamic factor model to extract the common component of the A-B flows.

tion aging, which lowered unemployment by about 0.6 ppt, and (ii) group-specific trends in participation, specifically a trend increase in participation for women up until the mid 90s, and a trend decline in participation for young workers since the early 80s, which lowered unemployment by about 1.6 ppt.

In contrast, a LFSS procedure identifies correctly neither the total effect of (i) and (ii), i.e., the total effect of demographic changes on US unemployment, nor the effect of (i), nor the effect of (ii).

First, a LFSS procedure appears to substantially *under*-estimate the total effect of demographics, as it finds that demographic forces lowered unemployment by only 1.2 ppt, substantially lower than our 2.2 ppt estimate. Using our stock-flow model, we explore the source of the bias suffered by the LFSS analysis, and we show that the two major changes in labor force participation (increased participation of women and decreased participation of young workers) account for the difference with our method. Since a LFSS approach does not take into account the mechanical link between the participation rate and the unemployment rate, it does not take into account how the two major trends in participation rates also affected the group unemployment rates, and it does not properly adjust for demographic changes.

Second, a LFSS procedure substantially *over*-estimates the contribution of population aging, i.e., the effect of (i) alone, to the trend in unemployment (with an estimate of -1.2 ppt versus a true contribution of -0.6 ppt). This result is important in the context of the literature on the determinants of secular unemployment movements.[6] That literature has relied on LFSS procedures to conclude that population aging has been the prime factor behind the trend in the unemployment rate over the past 50 years, and our results challenge this consensus.

The remainder of this paper is organized as follows. In the next section we present a stock-flow model of the labor market. This model is used in Section 3 to illustrate the limitations of the traditional LFSS approach. In Section 4 we present new methods for the adjustment of unemployment rates and Section 5 illustrates the methodology for the US labor market. Section 6 concludes and presents some directions for future research.

[6]See Perry (1970), Flaim (1979), Gordon (1998), Summers (1986), Shimer (1998), Shimer (2001).

2 A stock-flow model of the labor market

As mentioned in the introduction, the labor market is characterized by large flows of workers continuously taking place between employment (E), unemployment (U) and non-participation (N). This section presents an accounting framework that captures this dynamic nature of the labor market. The framework is a simple stock-flow model that allows for exogenous changes in the demographic structure of the population. We then show how the group-specific unemployment rates, the group-specific participation rates and the aggregate unemployment rate are functions of (i) the group-specific worker flows and (ii) the demographic composition of the population.

2.1 Group-specific unemployment and participation rates

We divide the population into I demographic groups. For our purpose, it will be useful to think about demographic groups as being defined by age and sex, although more general definitions are possible. At every instant, each worker can be in one of three states; employment (E), unemployment (U) and non-participation (N). We denote by U_{it}, E_{it}, and N_{it} the number of unemployed, employed and nonparticipants in demographic group $i \in \{1, .., I\}$ at time t. Further, let $LF_{it} = U_{it} + E_{it}$ be the size of the labor force in group i and $P_{it} = U_{it} + E_{it} + N_{it}$ the population size of group i.

The unemployment and participation rates of group i at time t are correspondingly defined as

$$\begin{cases} u_{it} = \frac{U_{it}}{LF_{it}} \\ \\ l_{it} = \frac{LF_{it}}{P_{it}} \end{cases} \tag{1}$$

To make the model consistent with data availability, we consider a continuous time environment in which data are available only at discrete dates. For $t \in 1, 2...T$, we refer to the interval $[t, t+1]$ as "period t". We assume that during period t, the instantaneous rates at which workers transition between labor market states are constant and given by $\{\lambda_{it}^{AB}\}$, where λ_{it}^{AB} is the hazard rate of an individual from demographic group i going from state $A \in \{E, U, N\}$ to state $B \in \{E, U, N\}$.

To model exogenous changes in the demographic structure of the population, we posit that the population of group i grows at an exogenous rate g_{it}, so that the population of group i at date t is given by

$$P_{it} = P_{it_0} e^{\int_{t_0}^{t} g_{i\tau} d\tau},$$

with P_{it_0} the population at some initial date t_0. While the population of type i individuals grows at a rate g_{it}, the numbers of unemployed, employed and nonparticipants need not, a priori, grow at the same rate. Specifically, denote g_{it}^U, g_{it}^E, and g_{it}^N the exogenous growth rates of U_{it}, E_{it}, and N_{it}. The population growth rate then satisfies $g_{it} = \frac{U_{it}}{P_{it}} g_{it}^U + \frac{E_{it}}{P_{it}} g_{it}^E + \frac{N_{it}}{P_{it}} g_{it}^N$.

During "period t", the fractions $\tilde{E}_{it} = \frac{E_{it}}{P_{it}}$, $\tilde{U}_{it} = \frac{U_{it}}{P_{it}}$ and $\tilde{N}_{it} = \frac{N_{it}}{P_{it}}$ of employed, unemployed and nonparticipants in demographic group i at instant $t + \tau$, $\tau \in [0, 1]$ follows the system of differential equations

$$\begin{cases} \frac{d\tilde{E}_{i,t+\tau}}{d\tau} = \lambda_{it}^{UE} \tilde{U}_{i,t+\tau} + \lambda_{it}^{NE} \tilde{N}_{i,t+\tau} - (\lambda_{it}^{EU} + \lambda_{it}^{EN}) \tilde{E}_{i,t+\tau} + \left(g_{it}^E - g_{it}\right) \tilde{E}_{it+\tau} \\ \frac{d\tilde{U}_{i,t+\tau}}{d\tau} = \lambda_{it}^{EU} \tilde{E}_{i,t+\tau} + \lambda_{it}^{NU} \tilde{N}_{i,t+\tau} - (\lambda_{it}^{UE} + \lambda_{it}^{UN}) \tilde{U}_{i,t+\tau} + \left(g_{it}^U - g_{it}\right) \tilde{U}_{it+\tau} \\ \tilde{N}_{i,t+\tau} = 1 - \tilde{E}_{i,t+\tau} - \tilde{U}_{i,t+\tau} \end{cases} . \tag{2}$$

The system (2) describes the flows of workers in and out of employment, unemployment and non-participation.

In practice, for the US and most OECD countries, population growth rates are small compared to worker transition rates, so that deviations from average growth rates ($g_{it}^E - g_{it}$ or $g_{it}^U - g_{it}$) are negligible compared to worker transition rates.[7]

Ignoring the influence of population growth and solving this system of linear first-order differential equations, it is easy to show that the group-specific unemployment rate u_{it} and the group-specific labor force participation rate l_{it} are functions of the present and past values of the six labor market transition rates $\left\{\lambda_{i,t-k}^{AB}\right\}$, $A, B \in \{E, U, N\}$ with $k \in [0, \infty)$.

[7]For instance, the US population increases at a rate of about 1% per year, so that $g_{it} \simeq 0.08\%$ at a monthly frequency, whereas the smallest monthly transition probability λ_{it}^{EN} fluctuates around 1% (Elsby et al. (2013)). Note that this assumption need not hold in countries with fast changes in population structure. In such cases, there exists a mechanical link between population growth rate and group-specific unemployment rate.

In particular, we can summarize the solution for u_{it} and l_{it} by

$$\begin{cases} u_{it} = \frac{U_{it}}{U_{it}+E_{it}} = u(\{\lambda^{AB}_{i,t-k}\}) \\ \\ l_{it} = \frac{U_{it}+E_{it}}{U_{it}+E_{it}+N_{it}} = l\left(\{\lambda^{AB}_{i,t-k}\}\right) \end{cases} \tag{3}$$

where $u(.)$ and $l(.)$ are functions of the worker transition rates $\{\lambda^{AB}_{i,t-k}\}$. Both functions are described in detail in Appendix A.

The important implication from (3) is that u_{it} and l_{it} are functions of the *same* underlying worker flows and thus cannot be treated as independent from each other. As we will see, this observation has important implications for the demographic adjustment of unemployment.

2.2 The aggregate unemployment rate

We combine the results from the previous section to express the aggregate unemployment rate as a function of (i) the underlying flows and (ii) the demographic composition of the population. The aggregate unemployment rate u_t is given by

$$u_t = \sum_{i=1}^{I} \omega_{it} u_{it} \tag{4}$$

where ω_{it}, the labor force share of group i is given by

$$\omega_{it} = \frac{LF_{it}}{LF_t}$$

with $LF_t = \sum\limits_{i=1}^{N} LF_{it}$.

To highlight how the labor force share depends on both the demographic structure of the population and the participation rates of each demographic group, we rewrite ω_{it} as a product of two terms

$$\omega_{it} = \frac{l_{it}}{l_t}\Omega_{it}, \tag{5}$$

where $l_t = \sum_{i=1}^{I} \Omega_{it} l_{it}$ is the aggregate labor force participation rate and where

$$\Omega_{it} = \frac{P_{it}}{P_t}$$

is the population share of group i.

When we combine (3), (4) and (5), the aggregate unemployment rate can be written as a function of (i) the transition rates $\left\{\lambda^{AB}_{i,t-k}\right\}$ and (ii) the population shares Ω_{it} with

$$u_t = \sum_{i=1}^{I} \omega_{it} u_{it}, \qquad \text{with} \qquad \begin{cases} \omega_{it} = \sum_{i=1}^{I} \frac{l\left(\lambda^{AB}_{i,t-k}\right)}{\sum_{j=1}^{I} \Omega_{jt} l\left(\lambda^{AB}_{j,t-k}\right).} \Omega_{it} \\ \\ u_{it} = u\left(\lambda^{AB}_{i,t-k}\right) \end{cases} \tag{6}$$

for $k \geq 0$ and $A, B \in \{E, U, N\}$.

Expression (6) shows that the labor force shares ω_{it} are not only functions of the population shares Ω_{it}, which capture the demographic structure of the population, but also of the underlying worker flows $\lambda^{AB}_{i,t-k}$. This point will be important below when we discuss the bias suffered by labor force share analyses.

3 Implications for labor force shift-share analyses

In this section, we describe the labor force shift-share procedure and show why it does not appropriately adjust the aggregate unemployment rate for demographic changes.

Recall from (6) that the aggregate unemployment rate is a weighted average of group-specific unemployment rates with weights given by the labor force share of each group. A labor force shift share analysis consists of constructing a counter-factual aggregate unemployment rate by holding labor force shares constant and letting only group-specific unemployment rates fluctuate. The goal of this procedure is to isolate the genuine fluctuations in unemployment from fluctuations generated by change in the composition of the labor force.

Specifically, a LFSS analysis starts from a first-order Taylor expansion of (4) with ω_{it} and u_{it} around their average value, so that the aggregate unemployment rate can be decomposed as

$$du_t = \sum_{i=1}^{I} u_i d\omega_{it} + \sum_{i=1}^{I} \omega_i du_{it}, \tag{7}$$

where $u_i = T^{-1} \sum_{t=1}^{T} u_{it}$ and $\omega_i = T^{-1} \sum_{t=1}^{T} \omega_{it}$ are respectively the average unemployment rate and the average labor force share of group i.[8]

[8]Alternatively, some authors linearize around the values of u_{it} and ω_{it} in some base year t_0. In practice,

From expansion (7), a LFSS method constructs a *demographic-adjusted* unemployment rate $\tilde{u}_t$ from

$$\tilde{u}_t \equiv \sum_{i=1}^{I} \omega_i u_{it}. \tag{8}$$

A LFSS-based analysis then proceeds to interpret changes in $\tilde{u}_t$ as capturing the "genuine" movements in unemployment coming from changes in group-specific unemployment rates. This counter-factual unemployment rate $\tilde{u}_t$ is then used to assess the state of the labor market and the business cycle.

However, from equation (6), we can see that both the labor force share ω_{it} and the group-specific unemployment rate u_{it} are functions of the transition rates $\left\{\lambda_{i,t-k}^{AB}\right\}$, $k \geq 0$ and $A, B \in \{E, U, N\}$. This implies that labor force shares and group-specific unemployment rates are *not* independent variables. Instead, they are *mechanically* linked, because both variables are driven by the same underlying worker flows taking place in the labor market. As a result, one cannot, as in a LFSS analysis, hold the labor force share fixed while letting the group-specific unemployment rate fluctuate (or vice-versa), and the counter-factual unemployment rates $\tilde{u}_t$ is ill-defined.

The direction and the magnitude of the bias suffered by $\tilde{u}_t$ depend on the six labor market flows and cannot be generally predicted. Indeed, each labor market flow generates a different mechanical link between unemployment and participation, so that the total co-movement between unemployment and participation generated by a demographic-specific trend, and therefore the bias suffered by $\tilde{u}_t$ (which does not take this effect into account), depends on the contribution of each flow to that demographic-specific trend. In our empirical application, we will highlight which flows are responsible for the major demographic trends in the US labor force, and we will discuss the direction of the bias suffered by a LFSS analysis of US unemployment.

---

the choice of base year makes little difference (see for example Katz & Krueger (1999)) and gives similar results to the ones obtained with a linearization around the mean.

4 A new approach for the demographic adjustment of unemployment

In this section we propose an alternative approach for the demographic adjustment of unemployment. The idea is to construct a counter-factual aggregate unemployment rate that is driven solely by common labor market fluctuations and not by demographic group-specific changes. To do so, we use a dynamic factor model of the worker flows to identify common variations across demographic groups and extract the common factors affecting the labor market.

4.1 Aggregate unemployment rate and worker flows

Our starting point is the same as in a traditional LFSS analysis like in Section 3, but we work directly with the worker flows instead of working with the stocks. Specifically, we start from equation (6), which define the aggregate unemployment rate as a function of (i) the population shares $\{\Omega_{it}\}$ for $i \in \{1,..,N\}$, and (ii) the worker transition rates $\left\{\lambda_{it-k}^{AB}\right\}$ for $i \in \{1,..,N\}$, $k \geq 0$ and $A, B \in \{E, U, N\}$.

Taking a first-order Taylor expansion of (6) with Ω_{it} and λ_{it-k}^{AB} around their average value gives

$$u_t = \sum_{i=\{1..I\}} \frac{l_i}{l} (u_i - u) \, d\Omega_{it} + \sum_{\substack{i=\{1..I\},\, k\geq 0 \\ A,B\in\{E,U,N\}}} \frac{\partial u_t}{\partial \lambda_{it-k}^{AB}} d\lambda_{it-k}^{AB}, \tag{9}$$

where the first component captures the contribution of changes in the demographic structure of the *population* (different from the labor force). Since the population shares $\Omega_{it} = \frac{P_{it}}{P_t}$ are only functions of the (exogenous) population growth rates g_{it}, we can identify the effect of changes in the population structure (like population aging) from movements in Ω_{it}.[9]

The second component of (9) captures the effect of changes in worker flows on the aggregate unemployment rate. To adjust the aggregate unemployment rate for demographic changes, we must decompose the different sources of movements in $\left\{\lambda_{i,t-k}^{AB}\right\}$. In particular, we must separate movements in $\left\{\lambda_{i,t-k}^{AB}\right\}$ that are group-specific (and thus attributable to demographics) from the movements in $\left\{\lambda_{i,t-k}^{AB}\right\}$ that are common across demographic groups

[9] We have $\Omega_{it} = \frac{P_{it}}{P_t} = \frac{P_{i0} e^{\int_0^t g_{i\tau} d\tau}}{\sum_{i=1}^{I} P_{i0} e^{\int_0^t g_{i\tau} d\tau}}$.

and thus not attributable to demographic-specific changes.

We will do so by means of a statistical model, in which the labor market flows are driven by two different disturbances: aggregate disturbances and demographic-specific disturbances. Using a dynamic factor model, aggregate disturbances will be identified as the changes in $\left\{\lambda_{i,t-k}^{AB}\right\}$ that are common across demographic groups.

4.2 Dynamic factor model

We now present a dynamic factor model that allows us to decompose the worker flows of the different demographic groups into a common component and several demographic-specific components.

Specifically, for each one of the six A-B transitions (i.e., from state A to state B), we have a panel of transition rates $\{\lambda_{it}^{AB}\}_{i=1,\ldots,I,t=1,\ldots,T}$ for the I demographic groups, and we use a dynamic factor model to extract the common component of the A-B transition rates.[10]

The dynamic factor model takes the same form for all six panels and includes a common factor structure and a separate stochastic trend for each demographic group (as in e.g., Otrok & Whiteman (1998)). Dropping the AB superscript for notational convenience, the behavior of the hazard rate λ_{it} is modeled as

$$\lambda_{it} = a_i f_t + \tau_{it} + \epsilon_{it}, \qquad \epsilon_{it} \sim IID(0, \sigma_{\epsilon,i}^2), \tag{10}$$

where a_i is the loading of demographic group i on the common factor f_t, τ_{it} is the trend that pertains to demographic group i and ϵ_{it} is the mean zero disturbance term that has variance $\sigma_{\epsilon,i}^2$. The model decomposes the hazard rates λ_{it} into a common component which is captured by f_t and weighted by a_i, and demographic-specific components captured by τ_{it} and ϵ_{it}. Notice that model (10) allows different demographic groups to be affected differently by the common factor.

[10]With $A, B \in \{E, U, N\}$, there are six different A-B transitions and thus six panels.

The common factor and trend components are modeled by

$$\begin{cases} f_t &= \phi f_{t-1} + \zeta_t, \qquad \zeta_t \sim IID(0, \sigma_\zeta^2) \\ \\ \tau_{it} &= \tau_{it-1} + \eta_{it}, \qquad \eta_{it} \sim IID(0, \sigma_{\eta,i}^2) \end{cases}, \tag{11}$$

where f_t follows an autoregressive process of order one with parameter $|\phi| < 1$ and disturbance term ζ_t. The demographic-specific trends τ_{it} follow random walk processes with disturbances η_{it}. The disturbances ζ_t and η_{it} have mean zero and variances σ_ζ^2 and $\sigma_{\eta,i}^2$, respectively. We assume that the disturbances ζ_t, η_{it} and ϵ_{it} are mutually uncorrelated across i and t. The latter assumption separates the common component from the group-specific components and allows us to separate common shocks from idiosyncratic shocks.[11]

The stationary specification for the common factor f_t reflects the idea that the shocks that are common across demographic groups are related to cyclical business cycle movements. The random walk specification for the idiosyncratic components reflects the idea that the demographic trends are slow-moving and persistent. We emphasize that the *only* assumption we need for separating the common and demographic-specific components is that the common shocks ζ_t are uncorrelated with the idiosyncratic shocks η_{it} and ϵ_{it}. The other specification choices are only guided by empirical considerations.[12]

The common component $a_i f_t$ is not identified without additional restrictions. In particular, we have $a_i^* f_t^* = a_i f_t$ when $a_i^* = a_i b$ and $f_t^* = f_t b^{-1}$, for all $b \in \mathbb{R}$. Therefore, we normalize the common factor by setting $\sigma_\zeta^2 = 1$.[13] The common factor is initialized by $f_1 \sim N(0, (1-\phi^2)^{-1})$ and the demographic specific trends are initialized diffuse by adopting $\tau_{i1} \sim N(0, 10^5)$.

The dynamic factor model can be written as a linear state space model, and the model parameters, collected in the vector $\psi = \{a_i, \sigma_{\epsilon,i}^2, \sigma_{\eta,i}^2, \phi, \}$, are estimated using maximum

[11]The assumption that the common and idiosyncratic components are uncorrelated is standard in the dynamic factor model literature, see for example Stock & Watson (2011).

[12]For example in the empirical application we also considered higher order autoregressive models for the common component as well as stochastic cycle specifications as in Harvey & Trimbur (2003). All these alternative specifications give very similar estimates for the common factors. Also, when we allow the demographic-specific components to follow autoregressive processes, the estimates for the autoregressive parameters indicate the presence of a unit root.

[13]Since, we are only interested in the product $a_i f_t$ this choice has no consequences and identical results are obtained by for example setting $a_1 = 1$.

likelihood, where the likelihood is based on the output of the Kalman filter, see Harvey (1989) and Durbin & Koopman (2012).[14]

Given the estimated parameters, denoted by $\hat{\psi}$, we can compute the minimum mean squared error estimates for the common factor and the demographic-specific trends by applying the Kalman filter smoother, as in Durbin & Koopman (2012, Chapter 4). The estimate for the common factor is denoted by $\hat{f}_t$ for $t = 1, \ldots, T$.

4.3 Constructing a demographic-adjusted unemployment rate

Using the minimum mean squared error estimate for the common factor $\hat{f}_t$ and the maximum likelihood estimates for the loadings $\hat{a}_i$ we construct counter-factual hazard rates $\hat{\lambda}_{it}$ for each demographic group using

$$\hat{\lambda}_{it} = \hat{a}_i \hat{f}_t. \tag{12}$$

These counter-factual hazard rates are based solely on shocks that are common across demographic groups and are not driven by group-specific changes.

After constructing these counter-factual hazard rates for each panel $\{\lambda_{it}^{AB}\}_{i=1,\ldots,I,t=1,\ldots,T}$ for $A, B \in \{E, U, N\}$, we can construct our demographic-adjusted unemployment rate from the definition of aggregate unemployment and equation (6):

$$\hat{u}_t = \sum_{i=1}^{I} \Omega_i \frac{l\left(\hat{\lambda}_{i,t-k}^{AB}\right)}{\sum_{j=1}^{I} \Omega_{jt} l\left(\hat{\lambda}_{j,t-k}^{AB}\right)} u\left(\hat{\lambda}_{i,t-k}^{AB}\right) \tag{13}$$

for $k \geq 0$ and $A, B \in \{E, U, N\}$ and where $u(.)$ and $l(.)$ are functions described in Appendix A.

As expression (13) makes clear, the demographic adjustment has two separate parts: (i) the population shares are held fixed at their average value Ω_i in order to get rid of group-specific changes in *population* growth rates, and (ii) the worker flows are only driven by the common factors of $\left\{\lambda_{i,t-k}^{AB}\right\}$, in order to get rid of group-specific demographic trends in the worker flows.

[14] Alternatively one could use standard Bayesian methods based on the Gibbs sampler, see for example Otrok & Whiteman (1998).

5 Demographic adjustment of US unemployment

In this section, we illustrate our approach with the demographic adjustment of US unemployment. After describing the major demographic changes that took place in the US labor market over the last 50 years, we present our demographic-adjusted unemployment rate, and we contrast our counter-factual rate with the rate implied by a LFSS procedure.

5.1 Demographic changes since 1960

Three major demographic changes took place in the US since 1960: (i) the aging of the population, (ii) the increase in the labor force participation rate of women between 1960 and the mid–90s, and (iii) the decline in the participation rate of young workers since the early 80s.[15]

The top-left panel of Figure 1 takes the starting point of any labor force share analysis and plots the labor shares of four demographic groups: younger than 25, prime age-male, prime age female and older than 55. We notice two major changes: first a marked increase in the labor force share of women up until the mid 90s, and second an uninterrupted decrease in the labor force share of young workers since the mid-70s. Since these two groups (especially young workers but also prime-age females to a lesser extent) have the highest unemployment rates (top-right panel), changes in their labor force shares will have the largest effects on the aggregate unemployment rate.

However, a lesson from this paper is that movements in labor force shares are a mix of exogenous changes in the population structure and changes in relative participation rates, with each mechanism possibly having a different effect on the unemployment rate. The bottom row of Figure 1 thus shows separately the population share and the labor force participation rates of the four demographic groups.

The major change in the population composition of the US is the aging of the baby boom generation. As we can see in the left-bottom panel, the aging of the baby boom generation led to successive "bumps" in the share of the different groups as the baby boom generation

[15] Another noticeable change is the rising participation rate of older workers since the early 90s. However, this change has had relatively little effect on the aggregate unemployment rate, because old workers' unemployment rate do not differ drastically from the average unemployment rate (in stark contrast with young workers).

goes through the different phases of life. Initially, there was an increase in the population share of young individuals (from the 1960s to the mid 70s), then an increase in the population share of prime-age individuals (from the early 80s to the mid-90s) and finally an increase in the share of older individuals since the late 90s.

In addition to population aging, the participation rates of the demographic groups also evolved over time, and different groups show different trends. First, there was a marked increase in the participation rate of women from the 60s until the mid 90s. Second, the participation rate of young workers has been on a downward trend since the 80s. Thus, the decline in the labor force share of the youngs since the mid-70s is due to two phenomena: (i) aging of the baby boom generation, and (ii) a declining participation rate since the early 80s.

5.2 A demographic-adjusted unemployment rate

We now construct a demographic-adjusted US unemployment rate using our proposed methodology, and we contrast this counter-factual unemployment rate with the rate implied by an LFSS analysis.

5.2.1 Data

The construction of our demographic-adjusted unemployment rate requires data on worker flows, and we measure workers' hazard rates using matched CPS micro data from February 1976 until December 2012. As described in detail in Appendix B, the hazard rates are corrected for the 1994 CPS redesign, for time-aggregation bias (Shimer (2012)) and for margin error (Poterba & Summers (1986)). We consider 11 demographic groups defined by gender (male or female) and age (16-24, 25-34, 35-44, 45-54, 55-64, 65+). We merge the 65+ age-cohort for males and females because of sample size.[16]

Figure 2 plots the average hazard rates of different demographic groups, smoothed and

[16] A natural question is the appropriate definition (i.e., size) of a demographic group. The level of disaggregation must be high enough to allow us to isolate demographic changes, but not too high to ensure a reasonable signal-to-noise ratio in the measured worker flows. By considering 11 groups (similarly to previous studies; e.g., Shimer (1998)), we strike a balance between these two criteria: the flows are still reasonably well measured and we can safely isolate the major demographic changes identified in the literature (e.g., Lazear & Spletzer (2012)). As a robustness check, we considered smaller (5-year) age bins groups and obtained similar results for the extracted common components.

standardized for clarity of exposition. A simple visual inspection reveals that some hazard rates co-move strongly, while others display little commonality. For instance, transitions between employment and unemployment are characterized by a strong common component, but transitions from employment to non-participation show markedly different demographic-specific trends. The dynamic factor model discussed in Section 4.2 is designed to extract the common component that underlies each type of transition.

5.2.2 Estimation

To extract the common component of each transition $A - B$ with $A, B \in \{E, U, N\}$, we proceed as described in Section 4.2. We estimate the parameters of the dynamic factor model specified by equations (10) and (11) using maximum likelihood. Given the parameter estimates we compute the minimum mean squared error estimates, or best linear predictors, for the common factors. Based on these and the estimated loadings we construct the demographic-adjusted hazard rates $\left\{\hat{\lambda}_{it}^{AB}\right\}$ from equation (12), i.e., the counter-factual hazard rates driven only by the common component.

Figure 3 shows the common factors identified for each type of transition. Some common factors are highly cyclical; in particular the Unemployment-to-Employment and the Employment-to-Unemployment common factors. Other common factors, such as the common factor for Employment-to-Nonparticipation transitions, display less, or no, cyclical movements. With the demographic-adjusted hazard rates in hand, we construct the demographic-adjusted aggregate unemployment rate based on equation (13).

5.2.3 A demographic-adjusted unemployment rate

Figure 4 shows the demographic-adjusted aggregate unemployment rate along with the original unemployment rate as well as the LFSS based demographic-adjusted unemployment rate as defined in (8).

Overall, our demographic-adjusted unemployment rate is highly correlated with the other two unemployment measures at cyclical frequencies, but it shows a much less marked downward trend. It has a lower level in the early 80s and a higher level in the latest recession.We now discuss the implications of our new demographic-adjustment of unemployment.

Assessing the state of the business cycle We first use our demographic-adjusted unemployment rate to contrast the different business cycles since the mid 70s.

Figure 4 shows that the three labor market peaks of 1979, 1989 and 2006 were remarkably similar with an unemployment rate bottoming at about 5 percent. Only the 2001 labor market peak stands out with a lower demographic-adjusted unemployment rate at about 4.2 percent. Interestingly, the exceptional level of unemployment at that time is in line with the perception of an exceptionally buoyant labor market linked to the high-tech boom and an exceptionally long expansionary phase.

Comparing troughs, the demographic-adjusted unemployment rate indicates that the latest recession was substantially worse than the early 80s recession with the labor market displaying a lot more slack: Without demographic changes, the unemployment rate would have been a full two percentage points higher in 2010 than in 1983. This result is in stark contrast with the unadjusted unemployment rate which suggests that the early 80s recession was the worse one (with a peak unemployment rate of 10.8 percent in December 1982 instead of 10 percent in October 2009), as well as with the LFSS-based adjusted unemployment rate, which points to similar amounts of slack in both recessions.

The contribution of demographics to the trend in unemployment To better appreciate the contribution of demographics to the trend in unemployment, Figure 5 plots the difference between our demographic-adjusted measure and the aggregate unemployment rate. We can see that demographic changes lowered unemployment substantially; by about 2.2 percentage points over the past 40 years.

As expression (13) made clear, the effects of demographics on movements in aggregate unemployment comes from two mechanisms: (i) changes in the population structure (changes in Ω_{it}), and (ii) group-specific trends in the worker flows.

To visualize the separate contributions of (i) and (ii), we also plot the contribution of (i) alone, i.e., the contribution of population aging only.[17] The aging of the baby boom lowered unemployment by about 0.6 ppt since the mid 70s, which implies that a large share of the

[17] As shown in (9), this is given by $u_t^{\Omega} \equiv \sum_{i=1}^{I} \frac{l_i}{l} (u_i - u) \Omega_{it}$. The effect of an exogenous change in the population size of a group on the aggregate unemployment rate will depend on (i) the participation rate of that group relative to the average aggregate participation rate ($\frac{l_i}{l}$) and (ii) the extent to which that group's unemployment rate differs from the average aggregate unemployment rate ($u_i - u$).

contribution of demographics comes from group-specific trends, in particular the increase in participation for women between 1960 and the mid–90s, and the decline in participation for young workers since the early 80s, as we will see in the next section.

Finally, the dashed line in figure 5 plots the contribution of demographics as estimated by an LFSS analysis.[18] We can see that a LFSS analysis *under*-estimates the total contribution of demographics and *over*-estimates the contribution of demographics coming solely from population aging (that is, coming solely from (i)). This latter over-estimation is important in the context of the literature on the determinants of secular unemployment movements (e.g., Shimer (1998)). Because that literature has relied on LFSS procedures to evaluate the contribution of aging to the trend in unemployment over the post war period, our results challenges the consensus that aging of the baby boom has been the prime factor behind the trend in unemployment over the past 50 years.

5.2.4 Explaining the difference with an LFSS adjustment

Relative to our demographic-adjusted measure, the LFSS-adjusted unemployment rate underestimates the downward effect of demographics on unemployment by about 1 percentage point since the mid 70s. We now argue that the difference owes to the two main demographic changes that were discussed in Section 5.1: (i) increasing participation rate of women up until the mid 90s, and (ii) trend decline in the participation rate of young workers since the early 80s.

First, some of the under-correction of the LFSS-adjusted unemployment rate comes from the rise in women participation in the first half of the sample. To see this, Figure 6 plots the actual and demographic-adjusted participation and unemployment rates of prime-age females.[19] Absent demographic-specific trends, females' unemployment rate would not have declined as much. In fact, the female unemployment rate would have been roughly 2 percentage points higher by 1995 (figure 6, bottom panel). Since a LFSS analysis completely ignores that channel, that 2 percentage points difference times the average labor force share of prime-age females (about 30 percent) explains that a LFSS-measure under-corrects the

[18] That is the difference $u_t - \tilde{u}_t$.

[19] The demographic-adjusted participation and unemployment rates of prime-age females are constructed using the counter-factual hazard rates $\left\{\hat{\lambda}^{AB}_{i,t-k}\right\}$ and the functions $u(.)$ and $l(.)$ (described in the Appendix) with $u_{it} = u(\hat{\lambda}^{AB}_{i,t-k})$ and $l_{it} = l(\hat{\lambda}^{AB}_{i,t-k})$.

effect of demographics by about 0.60 percentage points in the early part of the sample. The flow responsible for the trends in women participation and unemployment rates is the Employment-to-Nonparticipation transition rate –the E-N rate– depicted in Figure 8.[20] A decrease in λ^{EN} increases the participation rate, as women becomes less likely to leave the labor force, but it also lowers the unemployment rate, as employed workers are less likely to enter unemployment through a non-participation spell (as was first shown in Abraham & Shimer (2001)).

Turning to the trend decline in young workers' participation rate, Figure 7 shows that the young-specific trends in young workers' flows also lowered their unemployment rate by more than 2 percentage points since the early 80s (bottom panel, figure 7). A LFSS analysis (which ignores that effect) will thus under-correct the effect of young workers' declining participation rate on aggregate unemployment by about 2*0.2=0.4ppt (given a labor force share of young workers of about 0.2) over the sample period. The most important flows behind the decline in young workers' participation are the U-N and N-U flows (Figure 8), and the trend decline in the N-U rate and the trend increase in the U-N rate simultaneously lowered both participation and unemployment.

Taken together, these back of the envelope calculations show that the LFSS-adjustment under-corrects for the effect of demographics on unemployment by about 1 percentage point since the mid-70s, in line with the difference in trends observed in Figure 4.

6 Conclusion

This paper discusses new insights and methods for the demographic adjustment of the unemployment rate. By taking a dynamic view of the labor market, we show that the standard labor force shift share analysis cannot properly adjust the unemployment rate for the effects of demographic changes. The key insight is that the labor force share of a demographic group is mechanically linked to that group's unemployment rate, because both variables are driven by the same underlying worker flows. As a result, one cannot, as in a labor force shift-share analysis, use the labor force share to "control" for composition.

[20]This was first found by Abraham and Shimer (2001) and can be easily shown with a variance decomposition exercise as in Fujita & Ramey (2009). Details are available upon request.

We propose a new demographic-adjustment procedure that uses a dynamic factor model of the worker flows to separate aggregate labor market forces and demographic-specific trends.

Using the US labor market as an illustration, our demographic-adjusted unemployment rate indicates that the 2008-2009 recession was much more severe and generated substantially more slack than the early 80s recessions. Moreover, after demographic adjustment, the three labor market peaks of 1979, 1989 and 2006 are remarkably similar, and only the 2001 labor market peak appears different with an exceptionally low unemployment rate.

Our demographic adjustment also alters previous conclusions on the role of demographics for the trend in US unemployment since the mid 70s. In particular, we find that the contribution of the aging of the baby boom to unemployment's trend is much smaller than previously concluded with labor force shift share procedures.

References

Abraham, K. G. & Shimer, R. (2001), Changes in Unemployment Duration and Labor Force Attachment, NBER Working Papers 8513, National Bureau of Economic Research, Inc.

Barnichon, R. & Nekarda, C. (2012), 'The ins and outs of forecasting unemployment: Using labor force flows to forecast the labor market', *Brookings Papers on Economic Activity* .

Blanchard, O. J. & Diamond, P. (1990), 'The Cyclical Behavior of the Gross Flows of U.S. Workers', *Brookings Papers on Economic Activity* **2**, 85–155.

Del Negro, M. & Otrok, C. (2008), 'Dynamic factor models with time-varying parameters: measuring changes in international business cycles'. Working Paper.

Durbin, J. & Koopman, S. J. (2012), *Time Series Analysis by State Space Methods*, Oxford University Press, Oxford.

Elsby, M., Hobijn, B. & Sahin, A. (2010), 'The Labor Market in the Great Recession', *Brookings Papers on Economic Activity* **20**, 1–48.

Elsby, M., Hobijn, B. & Sahin, A. (2013), 'On the Importance of the Participation Margin for Labor Market Fluctuations'. Working Paper.

Engle, R. F. & Watson, M. W. (1981), 'A one factor multivariate time series model of metropolitan wage rates', *Journal of the American Statistical Association* **76**, 774–781.

Engle, R. F. & Watson, M. W. (1983), 'Alternative Algorithms for the Estimation of Dynamic Factor MIMIC and Varying Coefficient Models', *Journal of Econometrics* **23**, 385–400.

Flaim, P. O. (1979), 'The effect of demographic changes on the nations unemployment rate', *Monthly Labor Review* **102**, 13–23.

Fujita, S. & Ramey, G. (2009), 'The cyclicality of separation and job finding rates', *International Economic Review* **50**, 415–430.

Gordon, R. J. (1998), Inflation, flexible exchange rates, and the natural rate of unemployment, *in* M. N. Baily, ed., 'Workers, Jobs and Inflation', Brookings Institute, Washington, pp. 88–155.

Harvey, A. C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press, Cambridge.

Harvey, A. C. & Trimbur, T. M. (2003), 'General Model-Based Filters for Extracting Cycles and Trends in Economic Time Series', *Review of Economics and Statistics* **85**, 244255.

Hatzius, J. & Stehn, J. (2012), 'Comment on: The Ins and Outs of Forecasting Unemployment: Using Labor Force Flows to Forecast the Labor Market', *Brookings Papers on Economic Activity* **45**(2 (Fall)).

Herbertsson, T., Gylfi, Z. & Phelps, E. (2002), 'Demographics and Unemployment'. Institute of Economic Studies Working Paper No. W0109.

Horn, R. & Heap, P. (1999), 'The Age-Adjusted Unemployment Rate: An Alternative Measure', *Challenge* **42**, 110–115.

Kaitz, H. (1970), 'Analyzing the Length of Spells of Unemployment', *Monthly Labor Review* **93**, 11–20.

Katz, L. & Krueger, A. (1999), 'The High-Pressure U.S. Labor Market of the 1990s', *Brookings Papers on Economic Activity* **30**, 1–87.

Kose, M. A., Otrok, C. & Whiteman, C. H. (2003), 'International Business Cycles: World, Region, and Country-Specific Factors', *American Economic Review* **93**, 1216–1239.

Kose, M. A., Otrok, C. & Whiteman, C. H. (2008), 'Understanding the Evolution of World Business Cycles', *Journal of International Economics* **75**, 110–130.

Kudlyak, M., Reilly, D. & Slivinski, S. (2010), 'Comparing Labor Markets across Recessions'. Economic Brief, Federal Reserve Bank of Richmond, eb10-04.

Lazear, E. & Spletzer, J. (2012), 'The United States Labor Market: Status Quo or A New Normal?'. Working Paper.

Otrok, C. & Whiteman, C. H. (1998), 'Bayesian leading indicators: Measuring and predicting economic conditions in iowa', *International Economic Review* **39**, 997–1014.

Perry, G. L. (1970), 'Changing Labor Markets and Inflation', *Brookings Papers on Economic Activity* **1**, 411–448.

Perry, G. L. (1972), 'Unemployment Flows in the U.S. Labor Market', *Brookings Papers on Economic Activity* **3**, 245–278.

Polivka, A. & Miller, S. (1998), The CPS After the Redesign: Refocusing the Economic Lens, *in* J. Haltiwanger, M. Manser & R. Topel, eds, 'Labor Statistics and Measurement Issues', University of Chicago Press, Chicago.

Poterba, J. M. & Summers, L. H. (1986), 'Reporting Errors and Labor Market Dynamics', *Econometrica* **54**(6), 1319–38.

Quah, D. & Sargent, T. J. (1993), A Dynamic Index Model for Lareg Cross Sections, *in* J. H. Stock & M. W. Watson, eds, 'Business Cycles, Indicators and Forecasting', University of Chicago Press, Chicago, pp. 285–306.

Shimer, R. (1998), Why Is The U.S. Unemployment Rate So Much Lower?, *in* B. S. Bernanke & J. J. Rotemberg, eds, 'NBER Macroeconomics Annual', MIT Press, Cambrige, pp. 11–61.

Shimer, R. (2001), 'The Impact of Young Workers on the Aggregate Labor Market', *Quarterly Journal of Economics* **116**, 969–1008.

Shimer, R. (2012), 'Reassessing the Ins and Outs of Unemployment', *Review of Economic Dynamics* **15**, 127–148.

Solon, G., Michaels, R. & Elsby, M. W. L. (2009), 'The Ins and Outs of Cyclical Unemployment', *American Economic Journal: Macroeconomics* **1**(1), 84–110.

Stock, J. H. & Watson, M. W. (1989), 'A One-Factor Multivariate Time Series Model of Metropolitan Wage Rates', *Journal of Business and Economic Statistics* **76**, 774–781.

Stock, J. H. & Watson, M. W. (2011), Dynamic Factor Models, *in* M. P. Clements & D. F. Hendry, eds, 'Oxford Handbook of Economic Forecasting', Oxford University Press, Oxford.

Stock, J. H. & Watson, M. W. (2015), 'Core inflation and trend inflation'. Working paper, available from http://www.princeton.edu/ mwatson.

Summers, L. (1986), 'Why Is The Unemployment Rate So Very High Near Full Employment?', *Brookings Papers on Economic Activity* **2**, 339–383.

Valletta, R. & Hodges, J. (2005), 'Age and education effects on the unemployment rate'. FRBSF Economic Letter, Federal Reserve Bank of San Francisco.

Appendix A: Derivation of $u(.)$ and $l(.)$ functions

Recall that during "period t", the number of employed ($E_{i,t+\tau}$), unemployed ($U_{i,t+\tau}$) and nonparticipants ($N_{i,t+\tau}$) of demographic group i at instant $t+\tau$, $\tau \in [0,1]$ follows the system of differential equations

$$\begin{cases} \frac{dE_{i,t+\tau}}{d\tau} = \lambda_{it}^{UE} U_{i,t+\tau} + \lambda_{it}^{NE} N_{i,t+\tau} - (\lambda_{it}^{EU} + \lambda_{it}^{EN}) E_{i,t+\tau} \\ \frac{dU_{i,t+\tau}}{d\tau} = \lambda_{it}^{EU} E_{i,t+\tau} + \lambda_{it}^{NU} N_{i,t+\tau} - (\lambda_{it}^{UE} + \lambda_{it}^{UN}) U_{i,t+\tau} \\ N_{i,t+\tau} = 1 - E_{i,t+\tau} - U_{i,t+\tau} \end{cases},$$

where we ignore the negligible contribution of population growth and therefore omit the tildes on E, U and N.

In matrix form, this gives

$$\underbrace{\begin{pmatrix} \overset{\bullet}{E} \\ U \end{pmatrix}_{i,t+\tau}}_{\dot{Y}_{i,t+\tau}} = L_{i,t} \underbrace{\begin{pmatrix} E \\ U \end{pmatrix}_{i,t+t}}_{Y_{i,t+\tau}} + \underbrace{\begin{pmatrix} \lambda^{NE} \\ \lambda^{NU} \end{pmatrix}_{i,t}}_{g_{i,t}} \tag{14}$$

with

$$L_{i,t} = \begin{pmatrix} -\lambda^{EU} - \lambda^{EN} - \lambda^{NE} & \lambda^{UE} \\ \lambda^{EU} & -\lambda^{UE} - \lambda^{UN} - \lambda^{NU} \end{pmatrix}_{i,t}.$$

The steady-state of the system, Y_t^*, is given by

$$Y_{i,t}^* = -L_{i,t}^{-1} g_{i,t}, \tag{15}$$

so that (14) can be written as

$$\dot{Y}_{i,t+\tau} = L_{i,t} \left(Y_{i,t+\tau} - Y_{i,t}^* \right).$$

The solution is given by

$$Y_{i,t+\tau} = e^{-L_{it}\tau} Y_{i,t} + \left(I - e^{-L_{it}\tau} \right) Y_{i,t}^*, \qquad \tau > 0$$

with

$$e^{-L_{it}\tau} \equiv \begin{pmatrix} e^{-\beta_{i1}\tau} & 0 \\ 0 & e^{-\beta_{i2}\tau} \end{pmatrix}$$

with $-\beta_{i1}$ and $-\beta_{i2}$ the real and negative eigenvalues of $L_{i,t.}$[21]

Using the solution at $\tau = 1$ gives

$$\begin{pmatrix} E_{i,t+1} \\ U_{i,t+1} \end{pmatrix} = e^{-L_{it}} \begin{pmatrix} E_{i,t} \\ U_{i,t} \end{pmatrix} + \left(I - e^{-L_{it}}\right) \begin{pmatrix} E^*_{i,t} \\ U^*_{i,t} \end{pmatrix} \tag{16}$$

so that the numbers of employed and unemployment at time $t+1$ are functions of the period t hazard rates as well as time t employment and unemployment.

Expression (16) defines a recursive relation for the unemployment rate $u_{it} = \frac{U_{it}}{U_{it}+E_{it}}$, so that given initial conditions $U_{i,0}$ and $E_{i,0}$, expression (16) defines a function $u(.)$ such that $u_{it} = u(\{\lambda^{AB}_{i,t-k}\})$ with $A, B \in \{E, U, N\}$ and $k \in [0, \infty)$.

Proceeding similarly with the participation rate $l_{it} = \frac{U_{it}+E_{it}}{U_{it}+E_{it}+N_{it}}$, (16) defines a function $l(.)$ such that $l_{it} = l(\{\lambda^{AB}_{i,t-k}\})$.

Appendix B: Construction of worker transition rates

This appendix describes our procedure to construct time series for the six hazard rates of each demographic group. Since the procedure is identical for each group, we omit demographic subscript for clarity of exposition.

After matching the CPS micro files over consecutive months, we can construct monthly transition probabilities for the six flows. We then operate three corrections to these transition probabilities. First, we correct the transition probabilities for 1994 CPS redesign, then for time-aggregation bias following Shimer (2012) and Elsby et al. (2013), and finally we correct for margin error following Poterba & Summers (1986).[22]

As shown by Abraham and Shimer (2001), the 1994 redesign of the CPS (see e.g., Polivka & Miller (1998)) caused a discontinuity in some of the transition probabilities in the first month of 1994. To adjust the series for the redesign, we proceed as follows. We start

[21] β_{i1}, β_{i2} are real and positive because $tr(L_{it}) < 0$ and $det(L_{it}) > 0$.

[22] The correction for margin error restricts the estimates of worker flows to be consistent with the evolution of the corresponding labor market stocks.

from the monthly transition probabilities obtained from matched data for each demographic group, and we take the post-redesign transition probabilities as the correct ones. The goal is then to correct the pre-94 value for the redesign. To do so, we estimate a VAR with the six hazard transition probabilities in logs estimated over 1994m1-2010m12, and we use the model to back-cast the 1993m12 transition probabilities.[23] With these 1993m12 values in hand, we obtain corrected transition probabilities over 1976m2-1993m12 by adding to the original probability series the difference between the original value in 1993m12 and the inferred 1993m12 value.

By eliminating the jumps in the transition probabilities in 1993m12, we are assuming that the discontinuities were solely caused by the CPS redesign. Thus, the validity of our approach rests on the fact that 1994m1 was not a month with large "genuine" shocks to the transition probabilities. Reassuringly, looking at other dataset that were not affected by the CPS redesign shows indeed no major discontinuity in 1994m1. First, the unemployment exit rate and unemployment entry rate computed from unemployment duration data, which were not affected by the CPS redesign (Shimer (2012) and Solon, Michaels & Elsby (2009), show no major discontinuity in 1994m1.[24] Second, the employment-population ratio computed with data from the Census Employment Survey (which was unaffected by the CPS redesign) shows no evidence of any discontinuity in 1994m1 (Abraham and Shimer, 2001).

[23] The number of lags were chosen to maximize out-of-sample forecasting performances. A similar VAR is used in Barnichon & Nekarda (2012) to forecast the six flow rates.

[24] Specifically, Shimer (2012) and Solon et al. (2009) use data from the first and fifth rotation group, for which the unemployment duration measure (and thus their transition probability measures) was unaffected by the redesign.

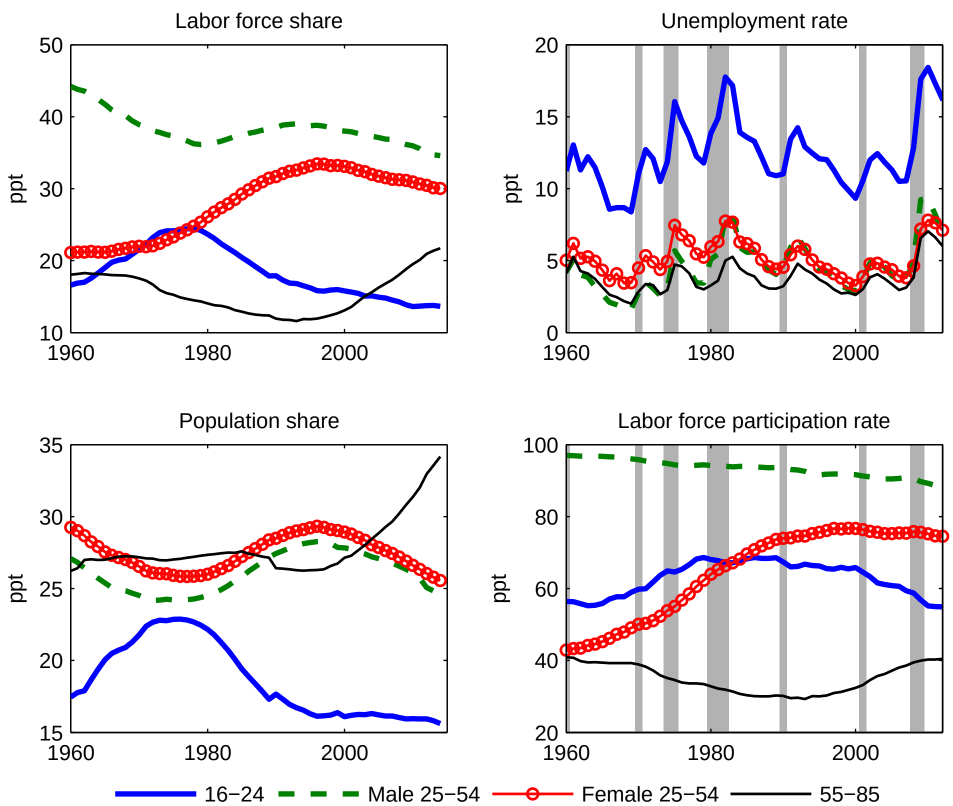

Figure 1: Demographic characteristics of the US labor market: labor force share (top left), unemployment rate (top right), population shares (bottom left) and labor force participation rate (bottom right) for four groups: Younger than 25, Male 25-54, Female 25-54, Older than 55. 1960-2012.

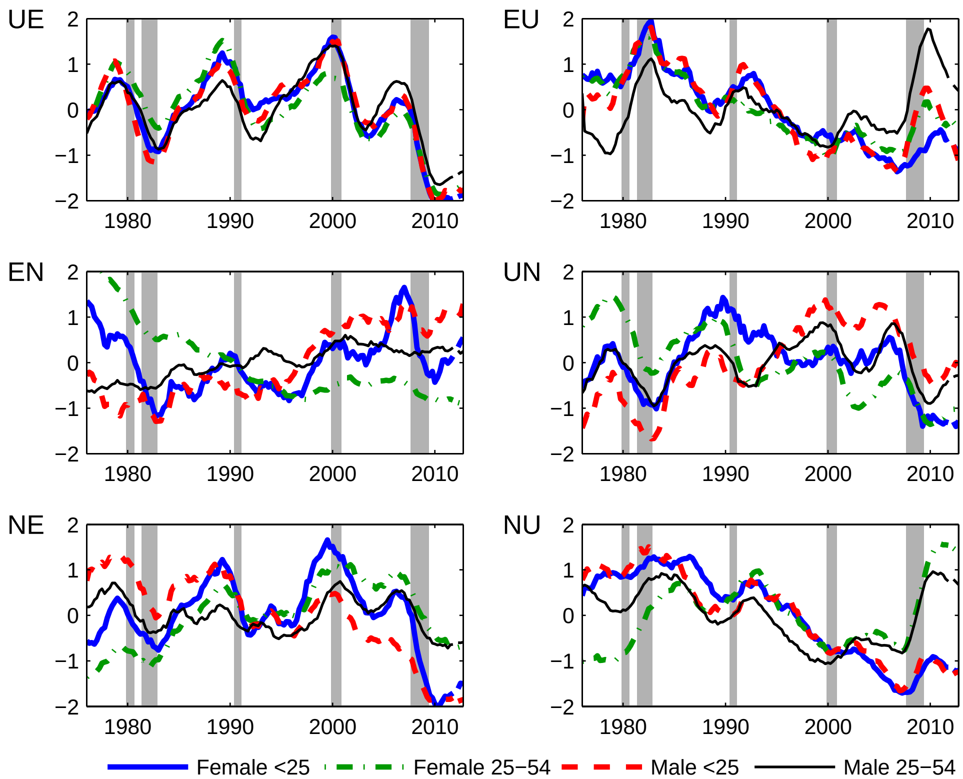

Figure 2: Smoothed and standardized transition rates of four demographic groups: young females, prime-age females, young males and prime-age males. 1976-2012.

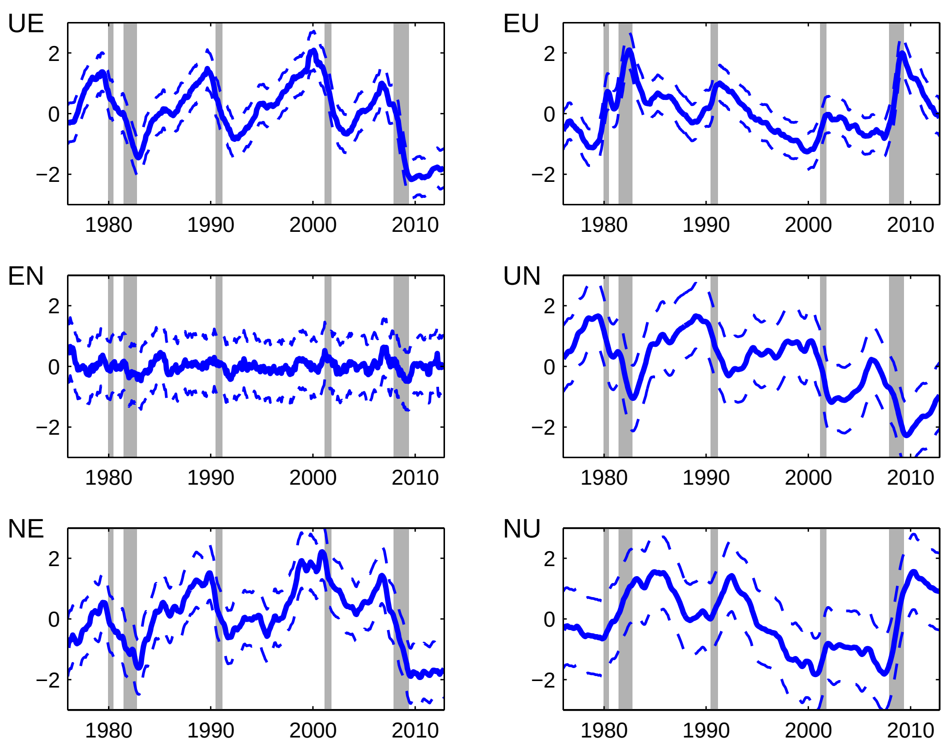

Figure 3: Common factors for the six hazard rates, 1976-2012. The dashed lines denote the 95 percent confidence interval.

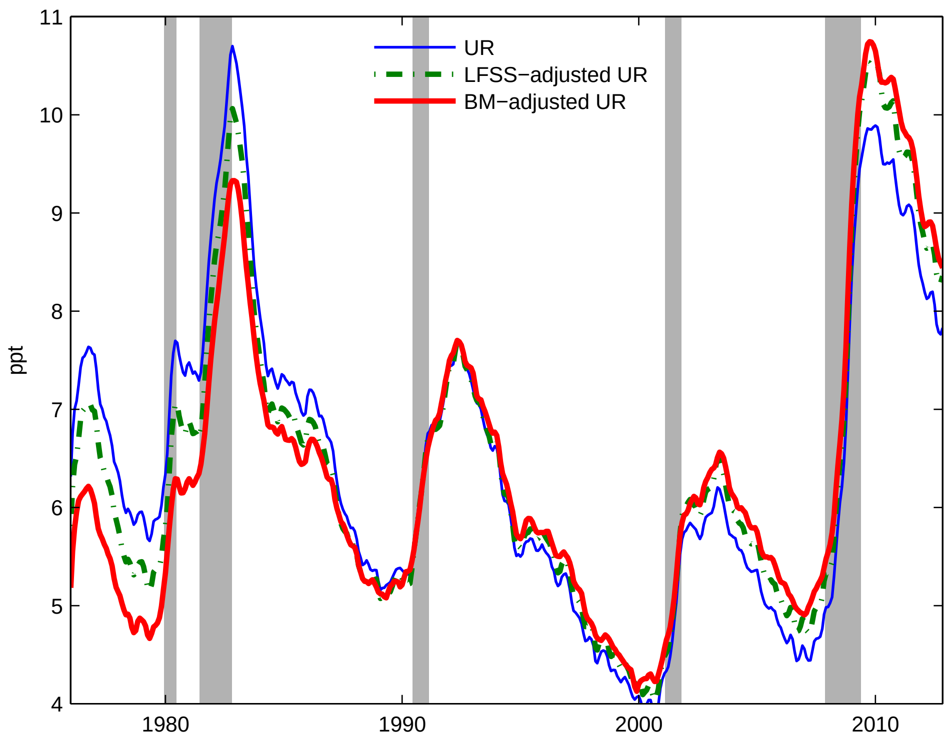

Figure 4: Demographic-adjusted unemployment rates. Unemployment rate (UR, solid line), LFSS-based demographic-adjusted UR (dashed line), Barnichon-Mesters(BM)-adjusted UR (thick line, based on the model of Section 4.2).

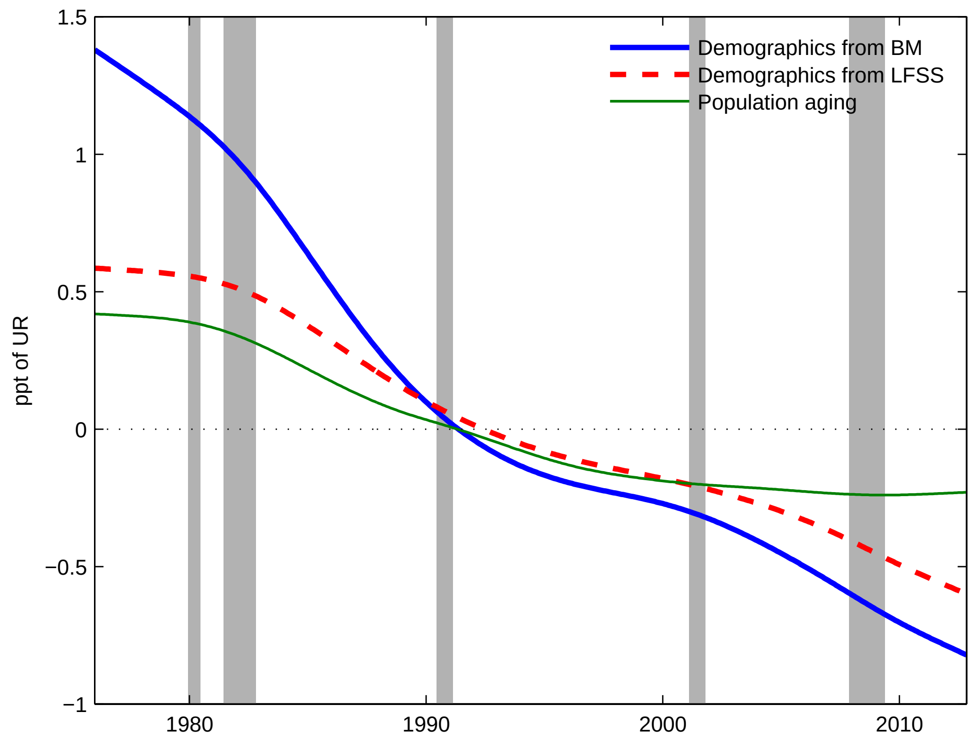

Figure 5: Contribution of demographics to unemployment (in ppt). Contribution evaluated using the Barnichon-Mesters (BM) procedure (thick line, based on the model of Section 4.2), contribution evaluated using the LFSS procedure (dashed line), and contribution of population aging (thin line).

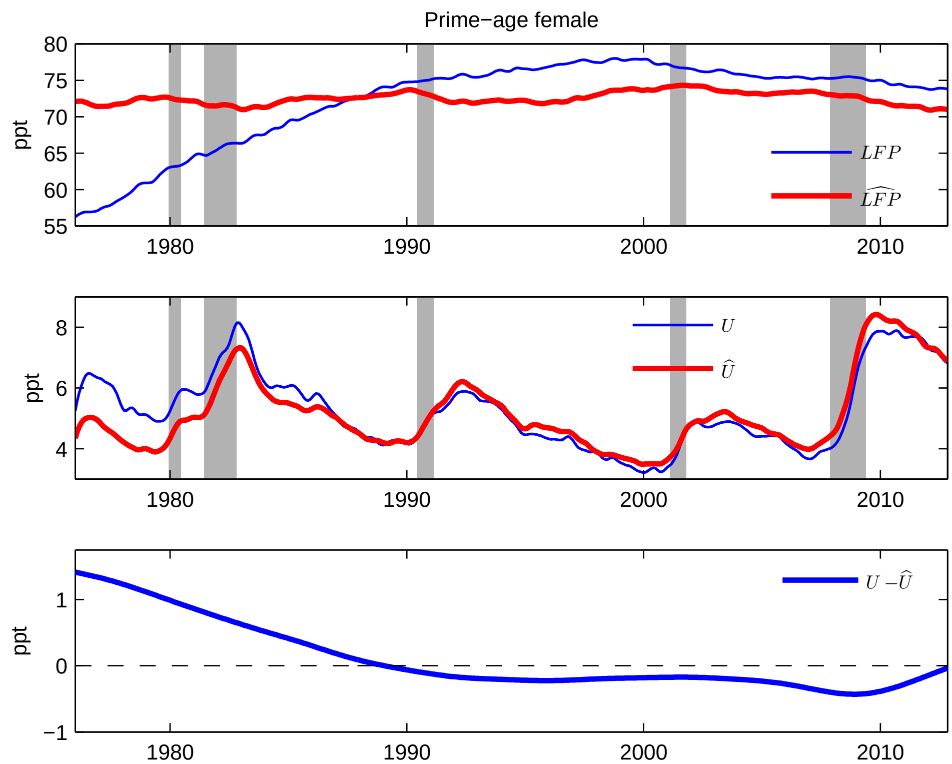

Figure 6: Demographic adjustment and prime-age females. Top panel: Labor force participation rates of prime-age female group (25-54); actual (thin line, labeled LFP) and counter-factual (thick line, labeled $\widehat{LFP}$). Middle panel: Unemployment rates; actual (thin line, labeled U) and counter-factual (thick line, labeled $\hat{U}$). Lower-Panel: Difference between U and $\hat{U}$. 1976-2012.

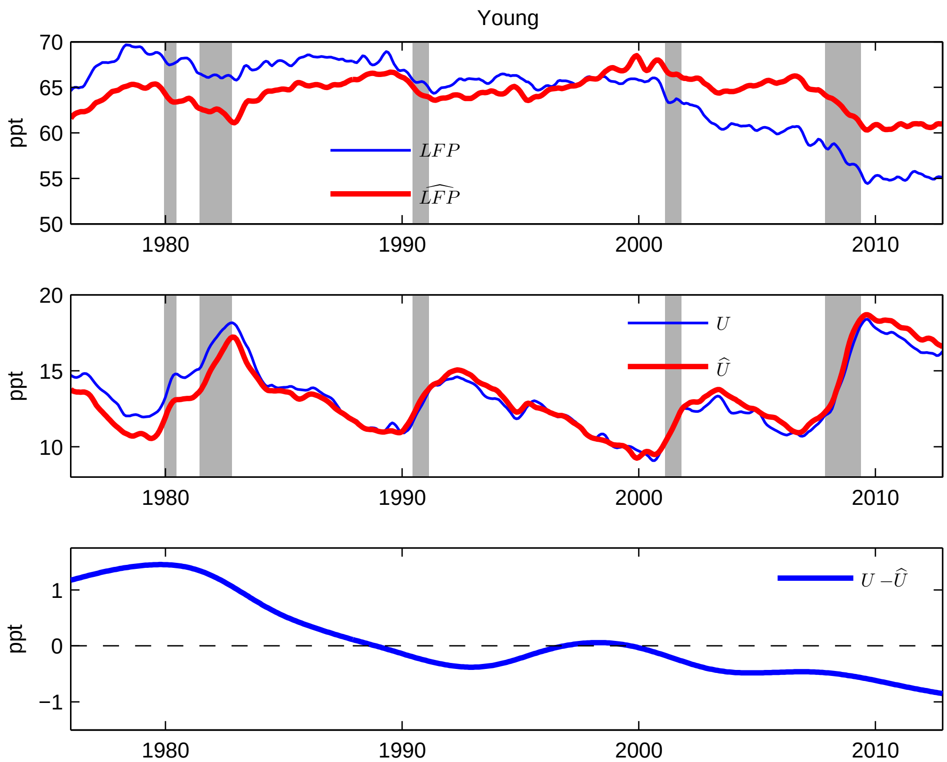

Figure 7: Demographic adjustment and young individuals (< 25). Top panel: Labor force participation rates of young individuals; actual (thin line, labeled LFP) and counter-factual (thick line, labeled $\widehat{LFP}$). Middle panel: Unemployment rates; actual (thin line, labeled U) and counter-factual (thick line, labeled $\hat{U}$). Lower-Panel: Difference between U and $\hat{U}$. 1976-2012.

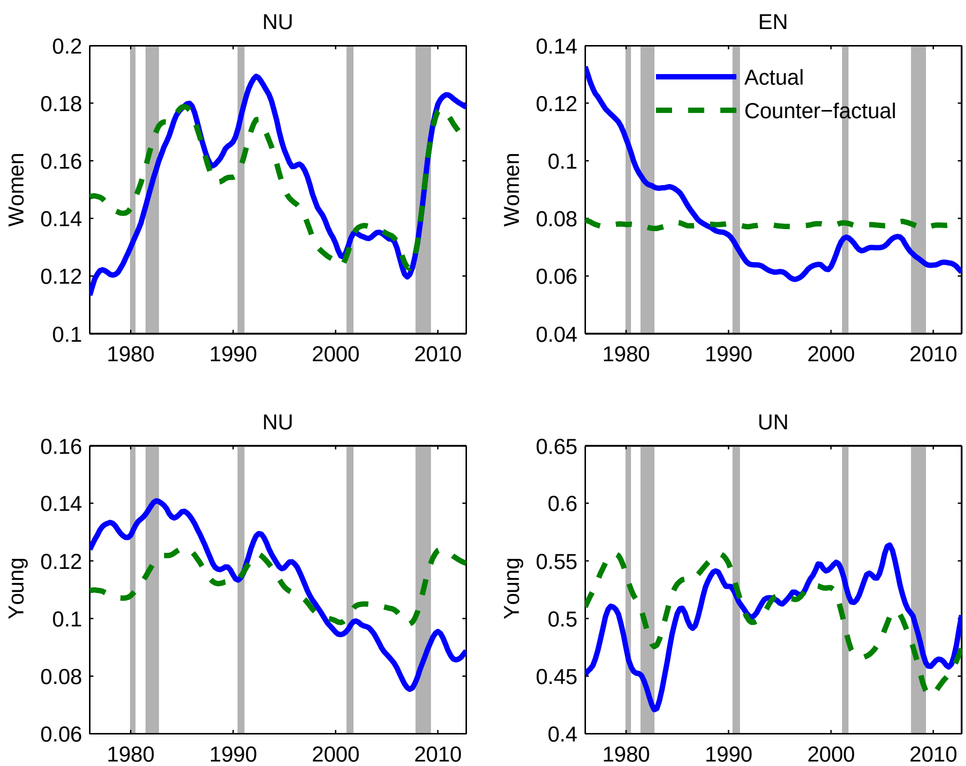

Figure 8: Selected transition rates for prime-age females (NU and EN) and young individuals (NU and UN). The thick lines denote the actual rates, and the dashed lines denote the counter-factual rates. 1976-2012. The series are smoothed with an MA(4).