Norges Bank

1 | 2024

Staff Memo

Employment trends in Norway

Authors:
Nicolai Ellingsen
Luca Fosso
Sigurd Mølster Galaasen

ISSN 1504-2596 (online)
ISBN 978-82-8379-307-9 (online)

Employment trends in Norway*

Nicolai Ellingsen †
Norges Bank

Luca Fosso ‡
European Central Bank

Sigurd Mølster Galaasen §
Norges Bank

January 17, 2024

Abstract

This paper outlines the recently developed method for assessing the trend level in employment rates adopted by Norges Bank. The approach employs a Bayesian VAR to decompose disaggregated employment data into trend and cyclical components, using quarterly labor market data on 30 demographic sub-groups. Applied to the time period 1984-2019, we show that the estimated trend picks up known historical factors contributing to slow-moving employment dynamics. Additionally, the cyclical employment component shows strong correlation with the output gap estimate for Norway.

1 Introduction

Ensuring that a large part of the population is employed is a fundamental objective of economic policy. To translate this goal into policy actions, it is necessary to both define what it means to have a high rate of employment and to measure how this objective changes over time. This is a challenging task, as the observed level of employment is the outcome of a wide range of factors that vary in terms of their observability, predictability, and duration. Nevertheless, to decide when and how policy should intervene it is essential to understand the nature of employment dynamics. In this paper we address this question by untangling the drivers of employment in Norway.

*We are grateful for valuable comments from Knut Are Aastveit, Karsten Gerdrup, Ørjan Robstad, Nicolò Maffei-Faccioli and Lorenzo Mori. The views expressed in the article are solely those of the authors and cannot be attributed to Norges Bank and to the European Central Bank.

†Norges Bank.: nicolai.ellingsen@norges-bank.no.

‡European Central Bank.: luca.fosso@ecb.europa.eu.

§Norges Bank.: sigurd-molster.galaasen@norges-bank.no.

Conceptually, it is useful for policy purposes to distinguish between trend and cyclical variation in employment. The trend component is interpreted as representing slow-moving structural factors, whereas the cycle component is associated with short-term economic fluctuations. Both components determine the current level of employment, but differs in terms of policy prescriptions. While the cycle component is primarily the concern of central banks and stabilization policies, the overall trend is potentially a matter for long-term labor market policies. Due to this division of policy, it necessary to assess in real time the relative importance of these factors.

There are several ways of performing such an assessment. One approach is by de-trending the aggregate time series using statistical filters such as for example the Hodrick and Prescott (1997)[1] method. Prior to the COVID-19 pandemic, however, Norges Bank adopted a more structural approach based on the assumption that the employment trend was driven entirely by demographic changes. In particular, the method involved fixing age-specific employment rates to their levels in a given reference year. Those levels were interpreted as reflecting their corresponding trend levels. In this setting, aggregate trend employment thus moves over time only because of changes in the age-composition of the population. Deviations from this trend are then attributed to cyclical variations. Importantly, the reference year was deemed to be neutral with respect to the cycle, meaning that the cycle component was assumed to be zero in that year.

In this paper we implement a methodology combining the two approaches. We first define demographic groups across a combination of age, sex and education levels. On the disaggregated data we then preform a trend-cycle decomposition using a Bayesian VAR in the spirit of Beveridge and Nelson (1981), allowing the trend to stochastically change over time.[2] Finally, we construct an aggregate trend as a size-weighted average of the group-specific trends. Thus, in contrast to the earlier Norges Bank approach, movements in the estimated aggregate trend arise from both demographic shifts and structural changes in group-specific trends.[3] Compared with simply de-trending aggregate employment directly, our bottom-up approach is better suited in settings characterized by large demographic changes and heterogeneous employment dynamics across popu-

1. See e.g. Veracierto (2008) for an application to the labor force participation and Krusell et al. (2017) on the employment rate. For alternative filtering methods, see e.g. Crump et al. (2019) and D'Amuri et al. (2021).

2. Our method is related to several recent contributions that allow for stochastic trends when decomposing macroeconomic data such as labor market and inflation dynamics (see e.g., Del Negro et al., 2017; Crump et al., 2019; Kamber and Wong, 2020; Ascari and Fosso, 2021; Hasenzagl et al., 2022).

3. An alternative which also allows for time varying trend movements is the approach adopted by the US Congressional Budget Office (Montes, 2018) based on Aaronson et al. (2006). In this approach the the employment rates of population sub-groups are estimated on age and cohort fixed effects, and various time-varying structural and cyclical co-variates.

lation sub-groups.[4] This characterization is particularly fitting for Norway, as shown in Bhuller and Eika (2020).[5]

We employ our method on Norwegian administrative data over the period 1984-2019, from which we identify the job status of individuals at the monthly level. Individuals are then allocated to one of 30 demographic groups, defined based on a combination of five age groups, three education levels and sex. Within each group, we aggregate the individual level data to create 30 quarterly group-specific time series for employment rates on which we perform a trend-cycle decomposition using the Bayesian VAR. Our estimated trends pick up known drivers of movements in the historical employment rate, such as changes in female labor participation and increased old-age participation. We also identify falling participation rate of low-educated workers. Moreover, the cyclical component resulting from the estimation seems to capture key business cycle movements and is highly correlated with Norges Banks' estimate of the output gap. We contrast our method with the previous approach adopted by Norges Bank, and show that the two methods provide different assessments of both the historical evolution and current assessment of the cyclical employment gap. In recent years, the difference arises primarily as a result of our method capturing an increased trend employment among older workers.[6]

The remainder of the article is organized as follows. In Section 2 we describe our data and how we measure employment at the individual and group level. In Section 3 we outline our statistical model, while Section 4 presents the decomposition estimates. We conclude by discussing the policy applications in Section 5.

2 Data

Data sources To create time series on employment rates disaggregated by age, education and sex we rely on Norwegian administrative data from Statistics Norway. The main data source is the employer-employee (EE) register providing us with start and end date of the near universe of wage contracts in Norway between 1984 and 2019. Individual's wage contracts are combined with background information on education and age using unique and anonymized personal identifiers.

4. Since Perry (1971) the bottom-up approach has been commonly adopted when studying labor market dynamics.

5. Bhuller and Eika (2020) decompose the decline in aggregate employment in Norway over the period 2000-2017, and show that accounting for both demographic changes and changes within demographic groups are important for explaining the aggregate decline.

6. This old-age employment growth has in other work been linked to the work-incentivizing features of the 2011 Norwegian pension reform, see e.g. Hernæs et al. (2016) and Galaasen and Kruse (2023).

Definition of employed A person is defined as employed in a given month if he or she has at least one active employment contract. If no employment contract is observed in any month of a given year, we cross check using the annual tax records to assess a person's employment status. An individual is then considered employed if (i) the annual salary income exceeds roughly USD 10,000 in 2023 nominal terms (the National Insurance basic amount) or (ii) net income above 1 basic amount from self-employment is reported that year.[7]

Demographic groups We split our sample into smaller units based on combinations of demographic characteristics. In particular we consider bins based on age, sex and education level. We consider five age groups, 16-24, 25-39, 40-54, 55-64 and 65-74 years of age, and three education groups, based on individuals' highest attained education level. The high-education group corresponds to a university degree, the medium-education group to a high school diploma, and the low-education group are those with less than high-school. A person's education characteristic is a fixed attribute, equal to the highest level of education observed for that person over the entire sample period.[8]. This will cause some right-censoring issues for the younger cohorts, as we typically do not know the final education status by the time our sample ends. Typically many in the youngest cohort after 2015 will be classified in the lowest education group. When estimating the model we therefore leave out the youngest cohorts during the later years in the sample to mitigate this problem.

Breaks in the data Benchmarking against the official aggregate employment series reveals that the coverage of employment relations in the EE register changes over time, progressively becoming more comprehensive toward the end of our sample period. However, our use of the tax data as an additional source for identifying employed persons almost entirely eliminates the discrepancy between our aggregate series and official statistics. Still, prior to 1993 we are unable to draw on the tax data to identify employment status among those missing from the EE register. This generates some discrepancy between our aggregate series and official statistics prior to 1993, as well as a jump in the disaggregate series from December 1992 to January 1993.

As our goal is to study trends in aggregate employment, we break-adjust our employment series such that our aggregate employment time series matches official statistics

7. As the detailed tax records start in 1993, prior to this year we cannot perform the last two steps.

8. For example, individuals who obtained a university degree in 2019 will be defined as highly educated over the entire sample period

each year. The break adjustment is performed by increasing employment in each demographic group by the same factor, thus keeping the relative shares across the groups constant. To illustrate, we have N demographic groups indexed by i with size and employment rates denoted as $(\mathrm{pop}_{it}, E_{it})$ in year t. We thus adjust the employment share with a factor $(1 + f_t)$ such that:

$$\sum_{i=1}^{N} \mathrm{pop}_{it} E_{it}(1 + f_t) = E_t^{\mathrm{off}} \tag{1}$$

where E_t^{off} is aggregate number of employed from official statistics. Thus, the break-adjusted employment in year t is given by: $B_t = \sum_i \mathrm{pop}_{it} E_{it} f_t$.

Table 1 shows the number of employees identified by each source in selected years. As we see, the break adjustment (B_t) is quite large for 1990.[9] This is mostly due to the fact that we do not observe detailed tax statements before 1993 which is used to cross-check the contractual employment definition. Our adjustment implicitly assumes that the missing employed B_t is distributed across the demographic groups according to their relative observed employment shares. If this assumption does not hold, our disaggregated employment series before 1993 might be biased as a result. However, the break in the series from December 1992 to January 1993 is quite similar across our demographic groups, which alleviates some of this concern.

Table 1: **Employees by identification method (selected years)**

Identified by	1990	2005	2015
Observed employment contract	1697	1975	2422
Self-employment income (tax statement) > 1 Basic amount	0	134	116
Wage income (tax statement) > 1 Basic amount	0	104	25
Break adjustment	338	29	1
Employment	2035	2242	2564

Table 1 displays the identification method for employment in our dataset for selected years using the method outlined in Section 2. Numbers in thousands of persons.

Employment rates Employment shares by age, sex and education levels are displayed in figure 1. Some well-known trends and facts stand out, illustrating substantial group-heterogeneity in the evolution of employment rates during the time period we look at. First, the large increase in female labor force participation is evident in the left panel,

9. Similar magnitudes are observed for all years prior to 1993

which shows a general increase in female employment rate from the 1980s until the early 2000s. Still, males are more likely to be employed throughout the sample period. Second, the middle panel shows that the core age groups, aged 25-54, have the highest employment rates. However, during the latter part of the sample period, the two oldest age groups experience growth in their employment rates that partially reduces the age gap. Looking at education in the right panel, we see a striking increase in the education-employment gap over time, driven by a substantial decline in the employment rate among individuals with low education.

Figure 1: **Employment rates by demographic groups**

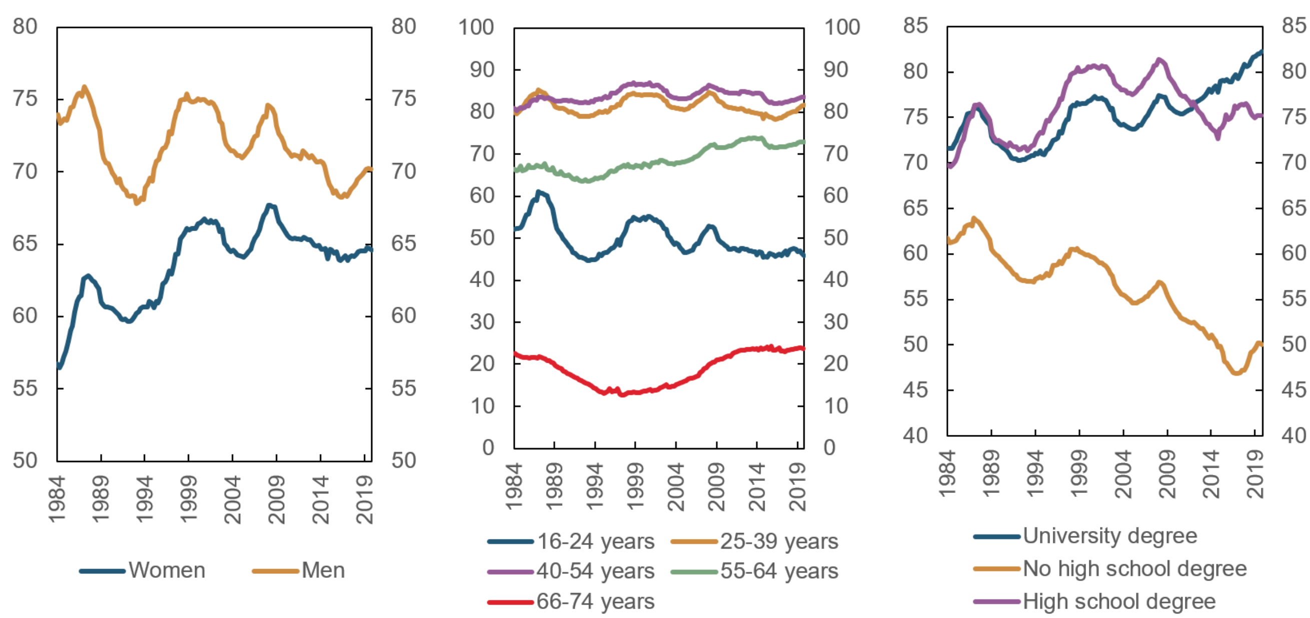

3 Empirical strategy

The ultimate goal of this empirical exercise is to obtain an estimate of the aggregate employment trend by exploiting the wealth of information embedded in our dataset. As mentioned above, the dataset can be described on the basis of three layers of disaggregation, from the less to the most disaggregated layer: (i) sex, (ii) age, and (iii) education. In order to get an estimate of the aggregate employment trend, we use a *bottom-up* approach and proceed as follows. First, we estimate the female (male) employment trends for each education level within all age groups using the empirical model presented be-low. This implies estimating a total of fifteen trends for both sexes[10], which we denote as $\bar{E}^{g}_{ij,t}$. The superscript $g = \{f, m\}$ refers to the sex, $i = \{16-24, 25-39, 40-54, 55-64, 65-74\}$ is the index of each age group and $j = \{\text{low}, \text{medium}, \text{high}\}$ is the index of each level of education. Second, using the population shares, we aggregate across education levels to obtain the employment trends for each age cohort

$$\bar{E}^{g}_{i,t} = \sum_{j} \frac{pop_{ij,t}}{pop_{i,t}} \bar{E}^{g}_{ij,t}$$

This allows to retrieve the sex-specific employment trends by age. Then, by aggregating over the five age groups, we calculate the female and male employment trends,

$$\bar{E}^{g}_{t} = \sum_{i} \frac{pop_{i,t}}{pop_{t}} \bar{E}^{g}_{i,t}$$

Finally, the aggregate employment trend is given by the weighted average of female and male trends. This bottom-up approach based on the use of population shares allows to have a clear understanding of the contribution to the slow-moving dynamics of the aggregate from the change of the composition of workers both in terms of education and age. Let us now discuss the technical details of the model the next paragraph.

3.1 Statistical model

Our empirical strategy builds on a linear state-space model that jointly characterises the secular and business cycle dynamics of the employment rates. Specifically, we assume that employment rates of the $N = 30$ demographic sub-groups evolves according to:

$$\underset{(N\times 1)}{E_t} = \underset{(N\times 1)}{\bar{E}_t} + \underset{(N\times 1)}{\hat{E}_t}, \tag{2}$$

$$\bar{E}_t = \bar{E}_{t-1} + u_t, \qquad u_t \sim N(0_N, \Sigma_u) \tag{3}$$

$$\Phi(L)\hat{E}_t = \varepsilon_t, \qquad \varepsilon_t \sim N(0_N, \Sigma_\varepsilon) \tag{4}$$

where E_t is an $(N \times 1)$ matrix with each element measuring the various sub-group's period t employment rate. The measurement Equation (2) decomposes data on employment rates into: (i) slow-moving trends $\bar{E}_t$; (ii) transitory (cyclical) components $\hat{E}_t$. The transition Equations (3) and (4) define a law of motion for the trends and the cycles. Trends

10. Three education levels within each of the five available age cohorts.

are assumed to be uncorrelated and to follow random walk processes, while the cycles follow a stationary reduced-form VAR process, where $\Phi(L) = I - \Phi_1 L - ... - \Phi_p L^p$ and Φ_k, for $k = 1, ..., p$, are the lag coefficients matrices of dimension $(N \times N)$. The vector stacking the permanent (Σ_u) and transitory (Σ_ε) innovations is assumed to be i.i.d. and distributed according to a multivariate normal distribution. In a nutshell, our model is equivalent to a VAR in deviations from its (constant) trends à la Villani (2009) with the notable difference that in this specification trends are allowed to be time-varying.

3.2 Estimation

The model is estimated separately for males and females, using Bayesian methods and employs the Kalman filter to extract the unobserved states in Equations 3 and 4. Therefore, we now discuss the priors of the model's parameters and initial conditions. The initial conditions of the unobserved states are assumed to be distributed according to:

$$\bar{E}_0 \sim \mathcal{N}\Big(\underline{E_0}, V_0\Big) \tag{5}$$

$$\hat{E}_0 \sim \mathcal{N}\Big(0, V(\Phi, \Sigma_\varepsilon)\Big), \tag{6}$$

where $\underline{E_0}$ is the pre-sample mean and V_0 is the (N x N) identity matrix. The initial condition of the cycles is a vector of zeros, as we assume that cycles symmetrically fluctuate around a zero mean. $V(\Phi, \Sigma_\varepsilon)$ is the unconditional variance of the initial conditions for cyclical components and it is always well defined as we impose stationarity on $\hat{E}_t$. Finally, the priors of the model's coefficients are distributed according to:

$$\Sigma_u \sim \mathcal{IW}(\kappa_u, (\kappa_u + N + 1)\underline{\Sigma_u}) \tag{7}$$

$$\Sigma_\varepsilon \sim \mathcal{IW}(\kappa_\varepsilon, (\kappa_\varepsilon + N + 1)\underline{\Sigma_\varepsilon}) \tag{8}$$

$$\text{vec}(\Phi)|\Sigma_\varepsilon \sim \mathcal{N}(\text{vec}(\underline{\Phi}), \Sigma_\varepsilon \otimes \underline{\Omega})\mathcal{I}(\text{vec}(\Phi)) \tag{9}$$

$\mathcal{IW}$ is the Inverse-Wishart distribution with κ degrees of freedom and mode equal to $\underline{\Sigma}$. We place a rather tight prior on the covariance matrix of permanent innovations $\Sigma_{\bar{u}}$ being diagonal by setting the hyperparameter $\kappa_{\bar{u}} = 100$. Notice, however, that this does not prevent the data to speak in favor of correlations among permanent innovations, if this is the case. Following Del Negro et al. (2017), the prior on the diagonal elements of $\Sigma_{\bar{u}}$ is conservative in limiting the amount of variance that is attributed to the secular trends – we scale the pre-sample data standard deviations by 100. Turning to the cyclical block, the priors for the lag coefficients are standard Minnesota priors with the overall tightness hyperparameter equal to 0.2, as suggested by Giannone et al. (2015), with the exception of the own-lag hyperparameter which is set equal to zero instead of 1, as we are characterizing the stationary behavior of the data. $\mathcal{I}(\text{vec}(\Phi))$ is an indicator function that is equal to value of one, if all the roots of $\Phi(L)$ are outside the unit circle, equal to zero, if the VAR is explosive. Since data are quarterly, the number of lags p in the stationary VAR is set equal to 4. Finally, we implement simulation smoothing techniques to generate the posterior distribution from the filtered measurements (see Carter and Kohn, 1994).

4 The estimated aggregate employment trend

Using the method outlined in Section 3 we obtain 30 estimated group-specific employment trends. The aggregate trend level at time t is then constructed as the sum the group trend levels, weighted by their corresponding population shares:

$$\bar{Y}_t = \sum_{i=1}^{30} S_{it} \Delta \bar{Y}_{it} \tag{10}$$

The evolution of the estimated aggregate employment trend over the period 1990-2019 is displayed in Figure 2, together with the historical employment rate in Norway. Overall, we estimate an increase in trend employment of about 3 percentage points in the early part of our sample period, and a slight decrease since the late 2000's. However, the reduction is much smaller compared with the actual employment rate decrease since after the 2008 financial crisis. Consequently, or method attributes the peak employment level to cyclical forces rather than changes in the underlying trend.

Figure 2: **Estimated aggregate employment trends**

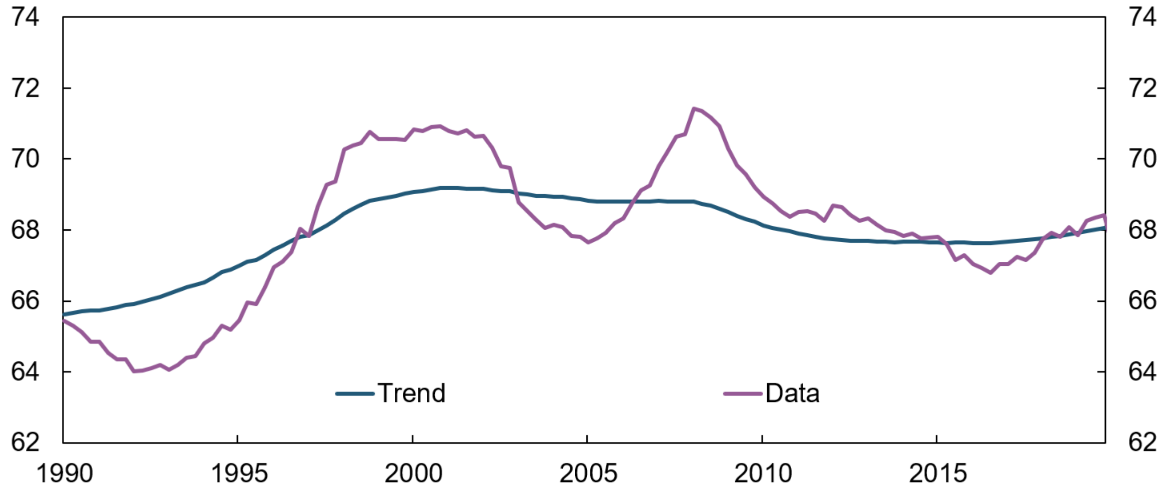

Figure 2 displays the historical aggregate employment rate and our estimated trend using the methodology in Section 3. The aggregate trend level in a given quarter-year is the weighted sum of the sub-groups' trend levels, using population shares as weights, following Equation (10).

The aggregate trend estimate in Figure 2 masks a large and heterogeneous evolution in the underlying group specific trends from 1990 until today. In Figure 3 we show employment trends at various level of sub-aggregation. Compared with females, the male employment trend has fallen sharply since the early 2000s (left panel). Older individuals have had a steady increase compared with younger individuals (middle panel), while individuals with low education have seen a large reduction.

Figure 3: **Disaggregated employment trends**

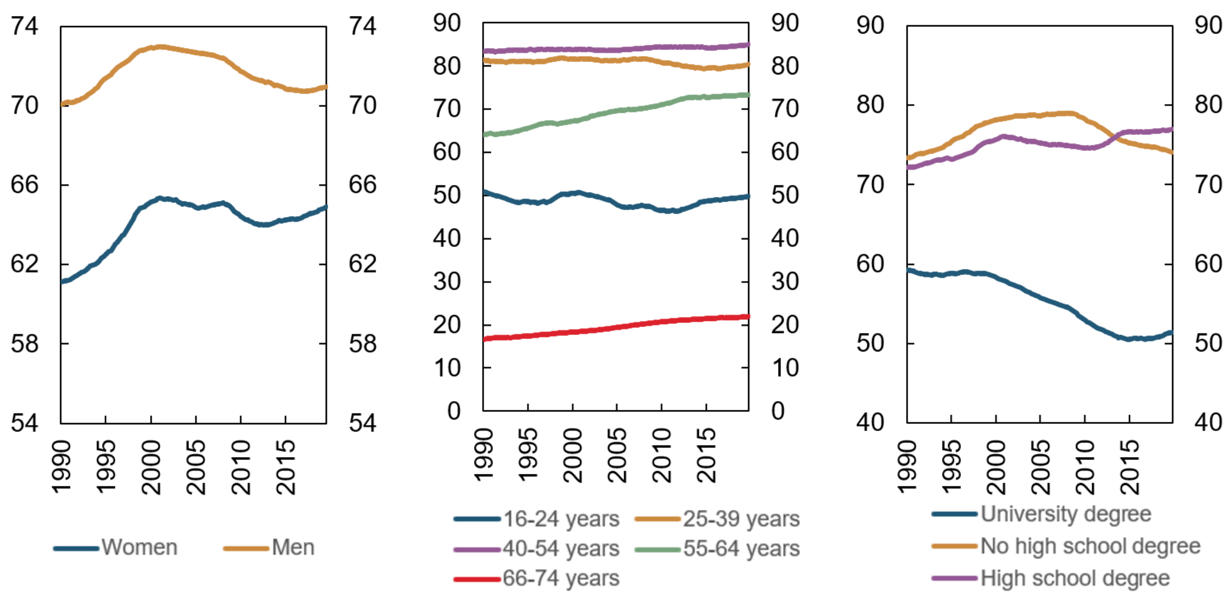

Figure 3 plots the estimated trend employment rate at various level of dis-aggregation. The left panel displays sex-specific trends. The middle panel displays age-specific trends. The right panel displays education-specific trends

Heterogeneous developments as those shown in Figure 3 pose a challenge for trend estimations based purely on demographic shifts. To illustrate the importance of accounting for changes in groups specific trends, we now perform a decomposition analysis of the change in the aggregate employment trend between 2013 and 2019. In particular, for each period $t > 2013Q1$ we break down the cumulative change in the aggregate trend into i) a pure demographic contribution, ii) a pure employment rate contribution, and iii) the combination between the two. Let $\Delta\bar{Y}_t = \bar{Y}_t - \bar{Y}_{t_0}$, with $t_0 = 2013Q1$ denote the cumulative change in employment trend, which we then decompose into

$$\Delta\bar{Y}_t = \underbrace{\sum_i S_{it_0}\Delta\bar{Y}_{it}}_{\text{Within}} + \underbrace{\sum_i \bar{Y}_{it_0}\Delta\bar{S}_{it}}_{\text{Between}} + \underbrace{\sum_i \Delta\bar{S}_{it}\Delta\bar{Y}_{it}}_{\text{Cross}} \tag{11}$$

The first term, "*Within*", measures the counterfactual change if we hold population shares S_{it} constant at their $t = t_0$ levels and vary only the group specific trend employment shares $\bar{Y}_{it}$. The second term, "*Between*", measures the counterfactual change if we fix employment shares to their $t = t_0$ levels and vary only the population shares. The last term, "*Cross*" accounts for the contribution of correlated changes in both employment and population. Importantly, the *Between* term measures the counterfactual trend level estimate we would have obtained if updating trend levels from 2013Q1 and onward using only information on demographic changes. The decomposition results are shown in Figure 4, together with the total change in the aggregate trend over the period 2013-2019. Ignoring changes in group specific trends would lead us to conclude with a cumulative decrease in the employment trend of -0.2 percentage point over the entire period. In contrast, our actual estimated trend grows by 0.4 percentage point over the same period. The force pulling the trend upwards is primarily an increase in the weighted average[11] group specific employment share, the *Within* term. However, we also see as small positive contribution from the *Cross* term. This means that there is a positive correlation between changes in trends and population shares; groups who become relatively larger in size (e.g. older age groups) also tend to experience an increase in trend employment over the period 2013-2019.

Figure 4: **Decomposing aggregate trend**

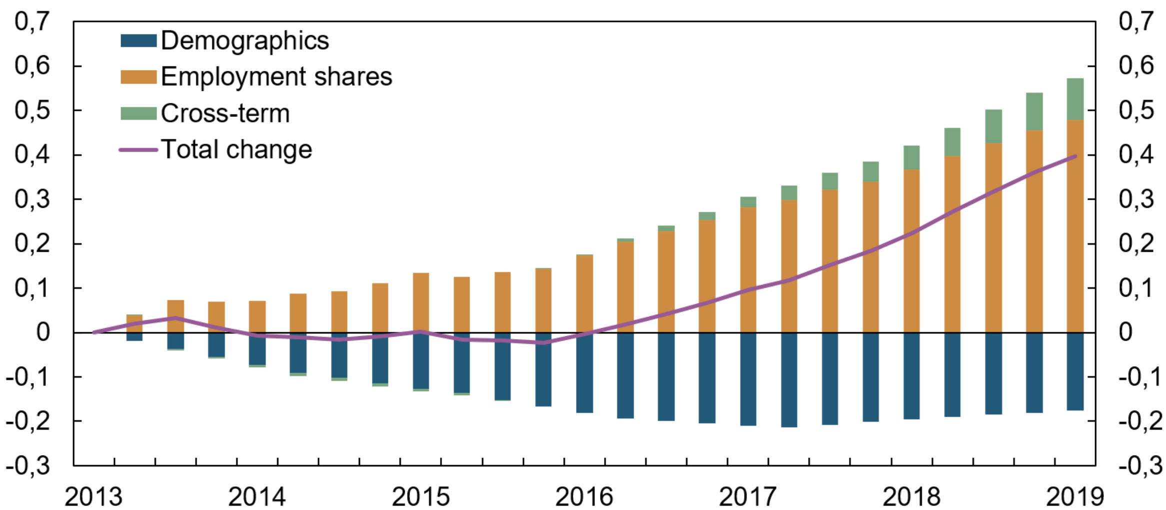

Figure 4 plots the decomposed cumulative change in the aggregate trend employment share over the period 2013-2019. The Within (Employment shares), Between (Demographics) and Cross-terms are obtained from the decomposition of the total change in Equation (11)

Finally, to assess whether the trend estimate in Figure 2 is in line with alternative activity measures, we now construct the employment gap as the difference between the observed aggregate employment share and its trend. As shown in Figure 5, the employment gap provides strikingly similar business cycle fluctuations as the more traditional output gap series.

11. Using $t = t0$ population shares as weights

Figure 5: **Employment gap**

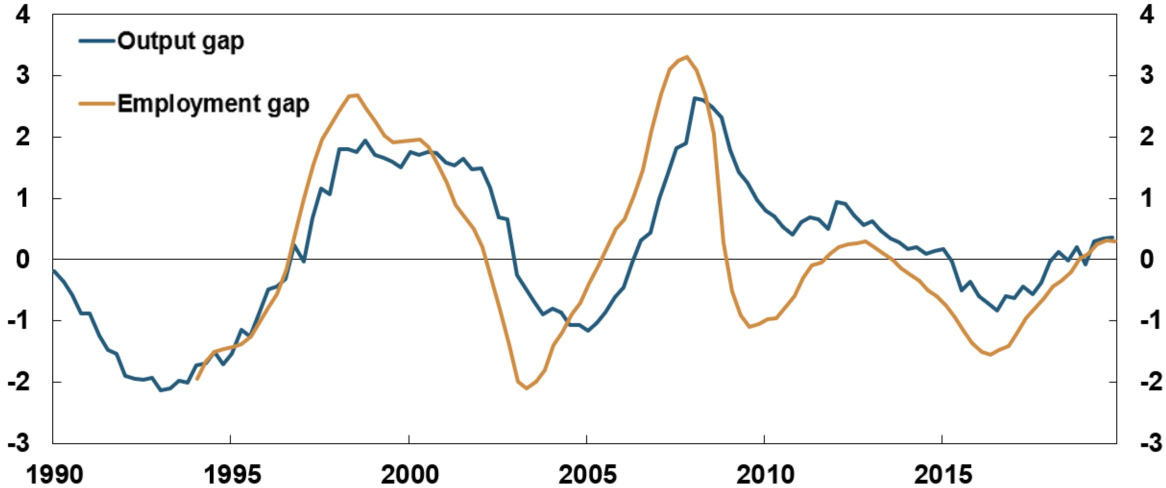

Figure 5 plots the output gap series for Norway and our employment gap series. The output gap is Norges Banks' official estimate, while the employment gap is measured as the difference between the actual employment rate and the trend employment rate estimated using the methodology in Section 3.

4.1 Comparison with the demographic-adjusted approach

The methodology developed in this paper provides a more flexible and robust trend assessment compared with the previous Norges Bank approach. The earlier methodology is essentially a demographic-adjusted approach, consisting of two steps. In the first step, the observed group specific employment rates are assumed to be on trend in a given base year. In the second step employment is assumed to evolve over time according to actual and projected population shares. This approach is restrictive as the estimate is sensitive to both the base year and the assumption of constant group trends.

As an example of this, we now contrast the trend estimate for the period 2007-2019 previously adopted by Norges Bank with our estimated trend. The comparison is shown in Figure 6. In the previous approach, indicated by the red line, the employment trend is derived by fixing employment within age groups to their observed 2013 levels, and then projecting the evolution using the observed aging of the population. In summary, our estimated trend, represented by the orange line, suggest a lower trend level in 2013 but a more positive trajectory, leading to a slightly higher employment trend in the last two years before the pandemic. The source of this difference is explained in Figure 4 where we decompose the change in the estimated aggregate trend estimates since 2013. The difference is primarily driven by increased trend employment among older workers, presumably caused by the 2011 Norwegian pension reform which stimulated labor supply among older workers (Hernæs et al., 2016).[12] The downward pressure on aggregate employment stemming from an aging population is partially offset by a rise in the estimated trend employment for older workers.

Figure 6: **Old and New Trend Estimate**

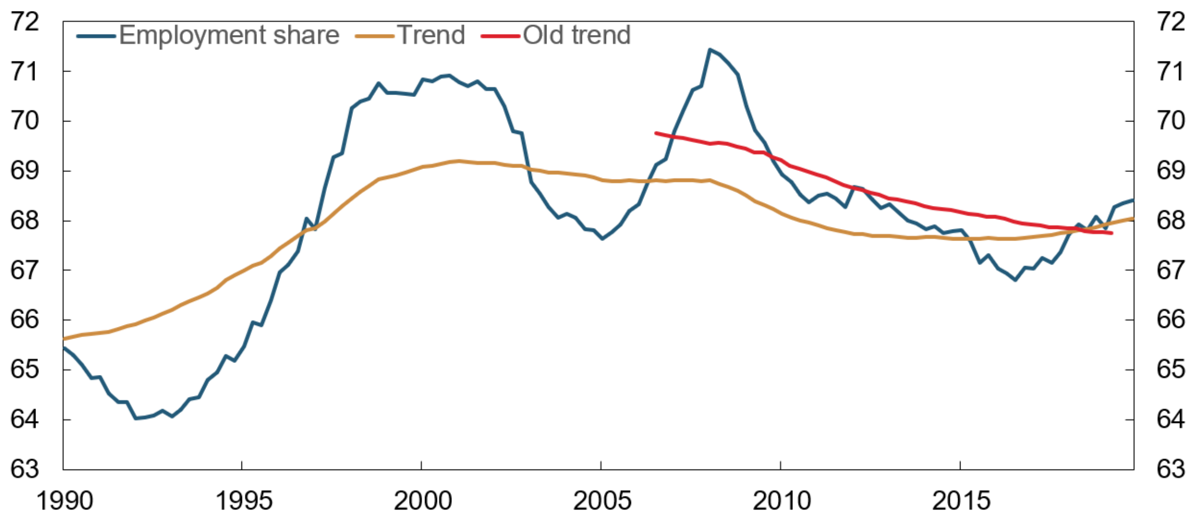

Figure 6 plots the old trend where employment within age groups are held fixed against our new estimated trend.

5 Policy application

This section explains how the method outlined in this paper is operationalized by Norges Bank's monetary policy department, both in analysing the current economic situation and in making forecasts. The estimated employment gap shown in Figure 5 is used when assessing the current temperature of the economy. The application to forecasting is less straightforward as additional elements are needed to forecast a trend level for actual employment. The basis for the forecast is individual trends for each 30 subgroups. We apply the random walk assumption from the model in Section 3 and assume that all trends remain constant from the latest observation in the data. Then we forecast the size of each group by the population forecast made by Statistics Norway. These population forecasts are not reported separately by education groups, so we add the assumption that education levels within age groups follow the trend from the last five years of reliable data (earlier than the end of the sample for younger cohorts). The forecasted

12. To account for the effect of the old age pension reform, the previous Norges Bank approach imposed a judgement-based increase in the trend employment of older workers over time.

group size and the estimated trend levels gives us a forecast on the level of domestic employment. For comparison with actual employment we add forecasted non-resident employment to the trend. Finally, the forecasted trend could be subject to judgement-based changes based on the assessment of other structural factors. For example, if our view on the NAIRU (non-accelerating inflation rate of unemployment) changes this could be reflected in our projected trend for employment.

6 Summary

The purpose of this paper has been to document the current method adopted by Norges Bank for estimating the trend level of aggregate employment. The method consists of a bottom-up approach, whereby we first estimate disaggregated trend levels for 30 demographic sub-groups using a Bayesian VAR, and the recover the aggregate trend as the population weighted average of the disaggregated trends.

Compared with simply de-trending aggregate employment directly, our bottom-up approach is better suited in settings characterized by large demographic changes and heterogeneous employment dynamics across population sub-groups. The reason is that it allows for both demographic changes and changes in trend employment within demographic groups to affect the aggregate trend level.

The method outlined in this paper replaces the previous Norges Bank approach for decomposing aggregate employment dynamics into cycle and trend components. We have illustrated how this the updated method improves and changes the assessment of the historical employment dynamics, and the contemporaneous trend level, by allowing group-specific trends to be time-varying.

Finally, we have evaluated how our estimated cyclical component co-moves with other measures of cyclical variation in activity. In particular, our employment gap (deviation from trend employment) shows strikingly similar dynamics since the 1990s to the more traditional output gap measure.

References

Aaronson, S., B. Fallick, A. Figura, J. F. Pingle, and W. L. Wascher. 2006. “The recent decline in the labor force participation rate and its implications for potential labor supply.” *Brookings Papers on Economic Activity* 2006 (1): 69–154.

Ascari, G., and L. Fosso. 2021. *The inflation rate disconnect puzzle: On the international component of trend inflation and the flattening of the Phillips curve.* Working Paper 733. De Nederlandsche Bank.

Beveridge, S., and C. R. Nelson. 1981. "A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'." *Journal of Monetary economics* 7 (2): 151–174.

Bhuller, M., and L. Eika. 2020. "Nedgang i sysselsettingen fra 2000–2017: –kan endringer i befolkningssammensetningen forklare alt?" *Søkelys på arbeidslivet* 37 (1-2): 20–37.

Carter, C. K., and R. Kohn. 1994. "On Gibbs Sampling for State Space Models." *Biometrika* 81 (3): 541–553.

Crump, R., M. Giannoni, S. Eusepi, and A. Sahin. 2019. "A Unified Approach to Measuring u." *Brookings Papers,* 143.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti. 2017. "Safety, liquidity, and the natural rate of interest." *Brookings Papers on Economic Activity* 2017 (1): 235–316.

Durbin, J., and S. J. Koopman. 2002. "A Simple and Efficient Simulation Smoother for State Space Time Series Analysis." *Biometrika* 89 (February).

D'Amuri, F., M. De Philippis, E. Guglielminetti, and S. Lo Bello. 2021. "Natural unemployment and activity rates: flow-based determinants and implications for price dynamics." *Bank of Italy Occasional Paper,* no. 599.

Galaasen, S. M., and H. Kruse. 2023. *Spousal Spillovers in the Labor Market: A Structural Assessment.* Working Paper 14. Norges Bank.

Giannone, D., M. Lenza, and G. Primiceri. 2015. "Prior Selection for Vector Autoregressions." *The Review of Economics and Statistics* 97 (2): 436–451.

Hasenzagl, T., F. Pellegrino, L. Reichlin, and G. Ricco. 2022. "A Model of the Fed's View on Inflation." *Review of Economics and Statistics* 104 (4): 686–704.

Hernæs, E., S. Markussen, J. Piggott, and K. Røed. 2016. "Pension reform and labor supply." *Journal of Public Economics* 142:39–55.

Hodrick, R. J., and E. C. Prescott. 1997. "Postwar US business cycles: an empirical investigation." *Journal of Money, credit, and Banking,* 1–16.

Kamber, G., and B. Wong. 2020. "Global factors and trend inflation." *Journal of International Economics* 122:103265.

Krusell, P., T. Mukoyama, R. Rogerson, and A. Şahin. 2017. "Gross worker flows over the business cycle." *American Economic Review* 107 (11): 3447–3476.

Montes, J. 2018. *CBO's Projection of Labor Force Participation Rate.* Working Paper 04. Congressional Budget Office.

Perry, G. L. 1971. "Labor force structure, potential output, and productivity." *Brookings Papers on Economic Activity* 1971 (3): 533–578.

Veracierto, M. 2008. "On the cyclical behavior of employment, unemployment and labor force participation." *Journal of Monetary Economics* 55 (6): 1143–1157.

Villani, M. 2009. "Steady-state priors for vector autoregressions." *Journal of Applied Econometrics* 24 (4): 630–650.

A Algorithm

The estimation is conducted following a Bayesian approach. The Gibbs sampler is structured according to the following steps:

1. Retrieve the distribution of latent states conditional on all the other parameter of the model

$$\bar{E}_{0:T}, \hat{E}_{-p+1:T} | \text{vec}(\Phi), \Sigma_u, \Sigma_\varepsilon, E_{1:T}$$

   Draws for the latent states can be obtained by using Durbin and Koopman (2002)'s simulation smoother. In addition, we also need to draws the initial condition $\bar{E}_0$ and $\hat{E}_{-p+1:0}$ in order to estimate the parameters in 3 and 4 in the next steps.

2. Draw the parameters of $\text{vec}(\Phi)$, Σ_u, Σ_ε conditional on the latent states

$$\text{vec}(\Phi), \Sigma_u, \Sigma_\varepsilon | \bar{E}_{0:T}, \hat{E}_{-p+1:T}, E_{1:T}.$$

   For given $\bar{E}_{0:T}$ and $\hat{E}_{-p+1:T}$, 3 and 4 are standard. In addition, we also already know the autoregressive matrices of the trend block in 3.

   The posterior distribution of the covariance matrix of permanent shocks Σ_u is given by

$$p(\Sigma_u | \bar{E}_{0:T}) = \mathcal{IW}(\underline{\Sigma}_u + S_u, \kappa_u + T),$$

where $S_u = \sum_{t=1}^{T}(\bar{E}_t - \bar{E}_{t-1})(\bar{y}_t - \bar{E}_{t-1})'$ is the empirical covariance matrix of permanent shocks to the trends. The posterior distribution of the stationary parameters in $\mathrm{vec}(\Phi)$ and Σ_ε is given by

$$p(\Sigma_\varepsilon|\hat{E}_{0:T}) = \mathcal{IW}(\underline{\Sigma_\varepsilon} + S_\varepsilon, \kappa_\varepsilon + T),$$

$$p(\mathrm{vec}(\Phi)|\Sigma_\varepsilon, \hat{E}_{0:T}) = \mathcal{N}\left(\mathrm{vec}(\hat{\Phi}), \Sigma_\varepsilon \otimes \left(\sum_{t=1}^{T}\hat{x}_t\hat{x}_t' + \underline{\Omega}^{-1}\right)^{-1}\right),$$

where $\hat{x}_t = (\hat{E}_{t-1}', ..., \hat{E}_{t-p}')'$ collects the VAR regressors,

$$\hat{\Phi} = \left(\sum_{t=1}^{T}\hat{x}_t\hat{x}_t' + \underline{\Omega}^{-1}\right)\left(\sum_{t=1}^{T}\hat{x}_t\hat{E}_t' + \underline{\Omega}^{-1}\underline{\Phi}\right), \quad S_\varepsilon = \sum_{t=1}^{T}\varepsilon_t\varepsilon_t' + (\hat{\Phi} - \underline{\Phi})'\underline{\Omega}^{-1}(\hat{\Phi} - \underline{\Phi}),$$

and $\varepsilon_t = \hat{E}_t - \hat{\Phi}'\hat{x}_t$ are the residuals of the stationary block of the model.