Immigration and Inequality: New Macroeconomic Evidence*

Francesco Furlanetto† Ørjan Robstad‡ Samad Sarferaz§

February 10, 2023

Abstract

In this paper, we reconsider the link between immigration and labor income inequality using detailed micro and macro data for Norway. Immigration has increased substantially in Norway during the last 20 years in response to several European Union enlargements to Eastern European countries. At the same time, several measures of income inequality have started rising, although not as abruptly as in other developed economies. Our analysis is feasible since Norway is one of the few countries for which detailed data on net immigration at the quarterly frequency are available together with micro-data on labor earnings collected by the tax authority covering the population of Norwegian workers at monthly frequency since 1997. We estimate a Factor Augmented Vector Autoregression (FAVAR) model tailored to include cross-sectional data to disentangle immigration shocks from other

*The authors thank Mikkel Riiser for impeccable research assistantship and Fabio Canova, Domenico Giannone and Giorgio Primiceri for useful comments. We thank also participants at ALUM/CEPR 2019 conference in Siracusa, CEF 2022 conference in Dallas, ASSA 2022 Virtual annual meetings, Dolomiti Macro meetings 2022 in Selva Val Gardena and various editions of the Norges Bank Spring Institute in Venastul. First presentation: June 2019. First draft: August 2022.

†Norges Bank and BI Norwegian Business School. Address: Bankplassen 2, PB 1179 Sentrum, 0107 Oslo, Norway. Email: francesco.furlanetto@norges-bank.no.

‡Norges Bank. Address: Bankplassen 2, PB 1179 Sentrum, 0107 Oslo, Norway. Email: orjan.robstad@norges-bank.no.

§KOF Swiss Economic Institute, ETH Zurich. Adress: Leonhardstrasse 21, 8092 Zurich, Switzerland. Email: sarferaz@kof.ethz.ch.

shocks driving the business cycle. Our main result is that a positive immigration shock increases inequality substantially with the effect being driven mainly by rich workers benefiting disproportionately from the increase in migration.

JEL classifications: C11; C32; E32

Keywords: immigration shocks, FAVAR model, inequality

1 Introduction

In this paper, we reconsider the effects of immigration on earnings using a combination of micro and macro data for Norway. The use of administrative data on earnings and detailed data on immigration flows to Norway at the quarterly level allows us not only to study the effects of immigration on average earnings, which are found to be relatively small in the literature (cf. Card (2009) among many others), but also to investigate the distributional effects of immigration and its impact on earnings inequality. The causal effect of immigration on labor income inequality is our key research question.

We believe that Norway is the ideal laboratory to study our research question for two reasons. First, Statistics Norway collects detailed data on immigration by covering the universe of migrants entering Norway with information on their reason for immigration, country of origin, gender, age, education level and sector of specialization. We are able to generate granular immigration series at the quarterly level and, importantly, we are able to reconstruct the aggregate net migration quarterly series from the detailed micro data as shown in Figure 1.[1] Second, the Tax authority collects data on individual earnings at the monthly level since 1997, thus allowing us to construct distributions of earnings data based on a cross-section of around 2.4 millions of observations at the end of the sample. These aggregate distributions can be further decomposed because we know several characteristics of each tax payer like gender, age, education, country of birth, sector of specialization (and even amount of financial wealth at annual frequency, a dimension which we currently do not exploit in the current paper). Our objective is to link the two data sets and extract the causal impact of immigration on earnings inequality at the macroeconomic level using time-series techniques.

[1] We used the quarterly net migration series to investigate the macroeconomic effects of job related migration shocks on the Norwegian economy in Furlanetto and Robstad (2019). Net migration data are usually available only at the annual level and the availability of a quarterly series was crucial to use Structural Vector Autogressive (SVAR) models in the context of the migration literature. We found that an increase in job related immigration lowers unemployment, lowers productivity and has no impact on house prices in Norway.

Figure 1: Immigration from official data and the micro data. Stock of immigrants on the left panel. Net immigration flow on the right panel

While the richness of data constitute the main reason to focus our interest on Norway, it is important to stress that immigration has been a key macroeconomic phenomenon over the last 20 years in the country. In fact, immigration has increased substantially in Norway in response to several European Union enlargements to Eastern European countries and constitutes the main driver of population growth as shown in Figure 3. The immigrant population increased from around 5 percent at the turn of the new millennium to 16 percent in 2016 with Poles, Lithuanians and Swedish constituting the largest communities. At the same time, several measures of income inequality have started rising, although not as abruptly as in other developed economies, see figure 2.

Figure 2: Gini index based on the earnings micro data for Norway

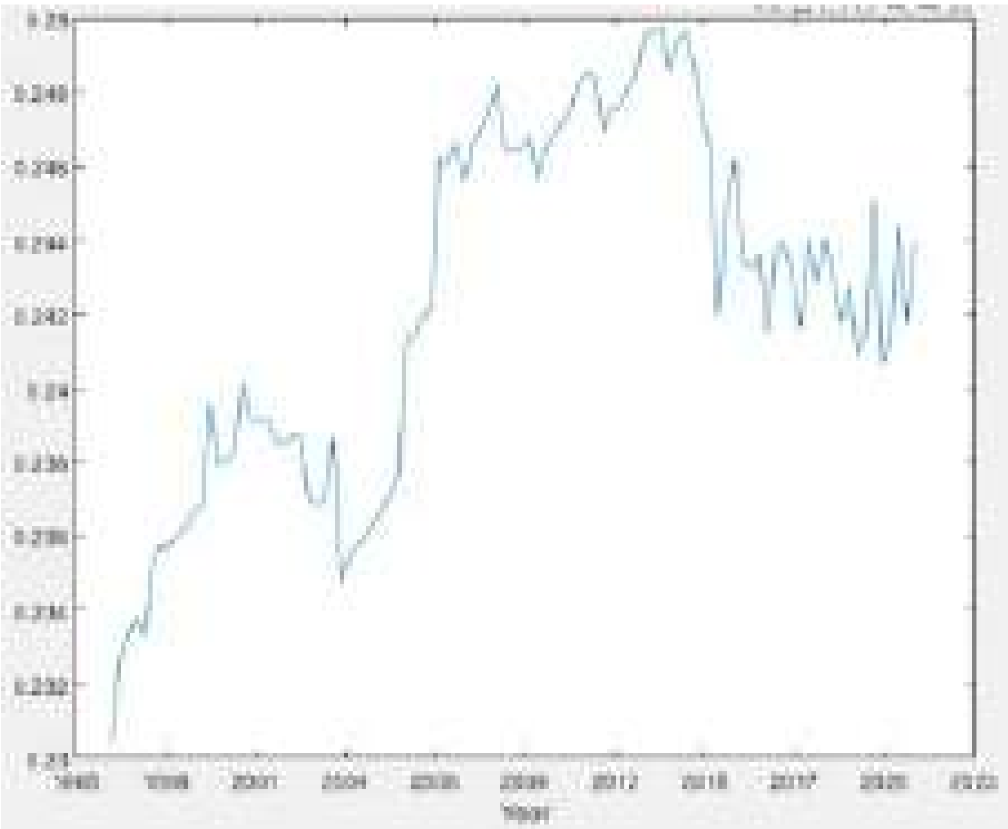

The link between immigration and inequality is not obvious. According to a standard supply-demand framework, immigrants are expected to lower the relative earnings of natives and previous migrants for whom they are close substitutes while wages and employment of complementary workers may even increase. On the other hand, migration may slow down the adoption of capital technologies and thus reduce the job polarization effects induced by technological changes (cf. Lewis (2011) and Basso, Peri and Rahman (2020) for US evidence and Furlanetto and Robstad (2019) for indirect evidence on Norway).

Figure 3: Population growth in Norway

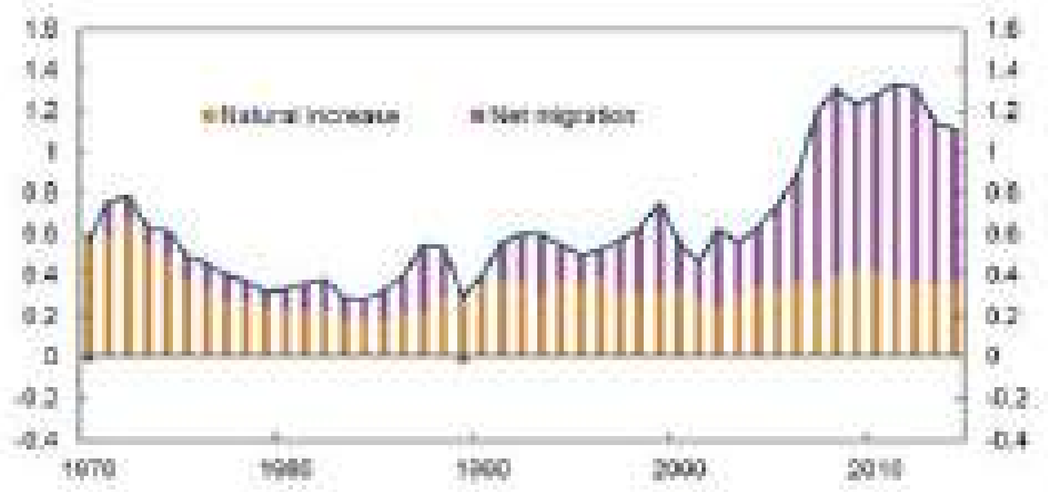

Our analysis is conducted at the level of the macroeconomy (although we will consider also a series of more disaggregated experiments) and thus we rely on techniques usually used to study the effects of other macroeconomic shocks, like monetary policy shocks. In practice, we combine a few aggregate time series, including the immigration series of interest in each experiment, together with the time series for the 100 percentiles summarizing the labor income distribution extracted from the micro data. Such a combination is performed in the context of a Factor Augmented Vector Autoregression (FAVAR) model, a model originally introduced by Bernanke, Boivin and Eliasz (2005) to trace the effects of monetary policy shocks over a large number (around 120) of macroeconomic time series. The variables are introduced in levels and macroeconomic shocks (including shocks to immigration) induce transitory deviations from the trend of the variables included in the system. Therefore, our focus is more on the cyclical component of labor income inequality rather than on its long-run drivers. Notably, the use of time series techniques (adapted to include cross-sectional data on earnings) distinguishes our approach from traditional analysis in the migration literature that are often conducted at the level of local labor markets using panel regressions. The adaptation of the FAVAR model to combine aggregate series with distributions derived from the micro data in a simple way constitutes the key contribution of this paper. An elegant (but more technically involved) alternative approach is provided by Chang and Schorfheide (2022) who use a functional VAR to study the effects of monetary policy shocks (conventional and informational) on the cross-sectional distribution of earnings and consumption (see also Chang, Gómez-Rodríguez and Hong (2022)).

As in the literature on monetary policy, a key aspect of the analysis is the identification strategy. In fact, we need to disentangle the exogenous component of immigration (immigration shocks) from the endogenous response of immigration to the state of the business cycle in Norway. Examples of immigration shocks that we have in mind are changes in regulation (like the various enlargements of the European Union that made possible for Eastern European workers to circulate and work freely within the continent) or variations in economic conditions in the source countries that are uncorrelated with the business cycle in Norway. In order to disentangle the exogenous and the endogenous component of immigration, we rely on a few restrictions on the sign and on the magnitude of the impulse responses to shocks in a re-interpretation of the max share approach originally proposed by Uhlig (2004) to identify the main driver of GNP. In practice, we identify two shocks, a general business cycle shock and an immigration shock, and we need some assumptions to the set the two shocks apart. Intuitively, we impose that variables related to the business cycle should be mainly (but not exclusively) driven by business cycle shocks while the immigration variable should be driven mainly (but not exclusively) by immigration shocks. The strategy is implemented by imposing a simple sign restriction over the ratio of employment over immigration. While both shocks are expected to move both variables in the same direction, it is reasonable to assume that a business cycle shock will move employment more than immigration while an immigration shock will move more immigration than total employment. Such a simple identification strategy has the advantage that no assumption on the behavior of aggregate wages and on the behavior of the earnings distribution is needed.

Aggregate immigration shocks have rather benign effects in our baseline model. They explain a non-minor share of fluctuations in employment and have limited effects on real wages. Most importantly, the labor income distribution responds rather little to the shock with the median percentile almost unaffected. However, not all percentiles respond in the same way. The highest percentiles of the distribution gain disproportionately from an immigration shock while the lowest percentiles exhibit only a slightly lower response than the median percentile. It is important to stress that, unlike in many micro studies, we are able to identify the *level* effect of immigration shocks and not only the differential effects on different sub-groups. While business cycle shocks are the main drivers of inequality in the short run, immigration shocks play a substantial role in the medium/long-run. Notably, the effect on inequality is substantially lower when we focus on the native population only. This is consistent with the standard result in the literature that previous immigrants are the most affected by an exogenous increase in immigration. Finally, an interesting result is that the effect on inequality is driven mainly by job related immigration. Family reunification and refugees have lower effects on inequality.

We contribute to two strands of the literature. Naturally, we integrate a recent and growing literature that studies the macroeconomic effects of immigration. Studies using using time series techniques include Kiguchi and Mountford (2019) for the US, Furlanetto and Robstad (2019) for Norway, Maffei-Faccioli and Vella (2021) for Germany, Schiman (2021) for Austria, d'Albis, Boubtane and Coulibaly (2019) and d'Albis, Boubtane and Coulibaly (2021) for several OECD countries. Studies using dynamic stochastic general equilibrium models include Smith and Thoenissen (2019), Hauser and Seneca (2022) and Olovsson, Walentin and Westermark (2021). In addition, we complement the literature on the effects of macroeconomic shocks on inequality. Coibion et al. (2017) study on the effects of monetary policy shocks on inequality using summary measures of inequality constructed from survey data while De Giorgi and Gambetti (2017) focus on the effects of technology and uncertainty shocks on consumption inequality. As far as we know, ours is the first paper to use macroeconomic techniques to investigate the effects of immigration on inequality.

Finally, we obviously relate to the enormous literature studying the effects of immigration at a more disaggregate level (see Ottaviano and Peri (2012), Card (2009) among many others and Kerr and Kerr (2011) for a survey). We see our approach as fully complementary. Contributions with a focus on Norway include Bratsberg and Raaum (2012), Bratsberg, Raaum and Røed (2014) and, more related to our research question, Hoen, Markussen and Røed (2022) who find that immigration from poor countries has steepened the social gradient in Norway (thus increasing labor income inequality) while immigration from rich countries has leveled the social gradient. The paper proceeds as follows. Section 2 describes the immigration and earnings data that we use in our analysis. Section 3 presents our empirical framework. Section 4 discusses our results. Section 5 presents some extensions to disentangle the effect of reason for immigration, education and country of origin. Finally, Section 6 concludes and presents some avenues for future research.

2 Data

The data used to construct the time series for the income percentiles and immigration is produced by combining several administrative registers from Statistics Norway. Income in these data is defined as labor income before taxes and transfers. The combined registers provide us with monthly information about the employment status and labor income for the majority of individuals in Norway from 1997:M1 to 2019:M12. In addition, we have information on demographic characteristics, including age, gender, education, occupation and date of immigration, reason for immigration and country of origin (if the person is an immigrant). An immigrant is defined as a person born in a country other than Norway. The data on labor income stems from the micro data in the employment registers at Statistics Norway. These registers covers close to the full population of employment contracts in our sample (around 2.4 millions of observations in the cross-section at the end of the sample).[2] In the baseline specification, we exclude individuals that work less

[2]Note that we only have register data for work contracts between employers and employees. Hence, self-employed workers are not part of the income distribution.

than 20 hours per week when we calculate the time series for the income distribution.[3] This is done to get a good representation of the labor income distribution for individuals whose main source of income comes from labor. The employment register has a break in 2015, when Statistics Norway started using a new higher quality register based on the monthly reporting of firms to the Norwegian Labour and Welfare Administration (NAV). In order to solve this issue and obtain a consistent time series, we use information from the annual tax returns. In the tax return data we have, among others, information about the annual labor income for the entire population in Norway. We adjust the level in the monthly income time series so that the annual growth rate in 2015 matches the labor income growth from the tax return for each income percentile.

The demographics attributes stem from the population registers. Since we do not have information about the education level of several migrants, we follow Jentoft (2014) to impute the education level. This imputation entails using other individual information, such as age, gender, length of residence, wage, occupation and country of origin to estimate the education level. The series included in our model refers to the stock of migrants present in Norway at each quarter. It is thus a slow-moving series that can be directly related with other stock variables like employment. Our focus on the stock of immigrants (and not on net migration) is inconsequential since the outflow of natives (emigration) is quantitatively small and relatively constant over time.

Other aggregate series, in addition to the immigration series, are combined with the earnings distribution in our baseline model. The remaining series include employment (measured in thousands of people, like the immigration series), the labor force participation rate and an average measure of real wages provided by Statistics Norway. All these variables exhibit a strong cyclicality and are therefore useful to identify a business cycle shock in our system. We prefer to use data on employment as a main indicator of the cycle since the data on GDP are heavily revised. In addition, employment and participation are somewhat more cyclical than unemployment in Norway, unlike in many

[3] We use the median income among all individuals included in each percentile when we produce the time series for the percentiles from the income distribution.

other countries.

3 Econometric Framework

We link immigration to labor income inequality using a slightly extended version of the FAVAR model described in Bernanke, Boivin and Eliasz (2005). Our model assumes that the co-movement of income percentiles can be explained by a few macroeconomic variables such as employment and immigration and some unobserved factors. We model the extracted factors and the observed macroeconomic variables jointly in a VAR, where the identification of immigration and business cycle shocks takes place via sign and magnitude restrictions. Based on the estimated relationship between the macroeconomic series and the income percentiles, we then trace the dynamic causal effects of the identified macroeconomic shocks on the jth income percentiles. In the following, we provide a more detailed description of this procedure.

Define $y_{j,t}$ as the jth income percentiles at time t for $j = 1, ..., 99$[4] and $t = 1, ..., T$ and consider the following model

$$y_{j,t} = \boldsymbol{\beta}_j^x \boldsymbol{x}_t + \boldsymbol{\beta}_j^\kappa \boldsymbol{\kappa}_t + \varepsilon_{j,t} \tag{1}$$

where $\boldsymbol{x}_t$ is a $m \times 1$ vector of observed macroeconomic variables, $\boldsymbol{\beta}_j^x$ is a $1 \times m$ coefficient vector, $\boldsymbol{\kappa}_t$ is a $k \times 1$ vector of unobserved factors, $\boldsymbol{\beta}_j^\kappa$ is a $1 \times k$ vector of factor loadings and $\varepsilon_{j,t}$ is an i.i.d. normally distributed error term with zero mean and variance σ_j^2. The idea behind Equation (1) is that the macroeconomic variables and the factors capture the common dynamics of the labor income percentiles and the error terms pick up their idiosyncratic movements.

As mentioned above, we assume that $\boldsymbol{x}_t$ and $\boldsymbol{\kappa}_t$ follow a joint VAR process

[4] Percentile number 100 has been removed from the system because it was creating instability in the estimates.

$$\boldsymbol{Z}_t = \boldsymbol{c} + \sum_{l=1}^{p} \boldsymbol{B}_l \boldsymbol{Z}_{t-l} + \boldsymbol{\nu}_t \tag{2}$$

where $\boldsymbol{Z}_t = [\boldsymbol{\kappa}_t' \ \boldsymbol{x}_t']'$, $\boldsymbol{c}$ is $n \times 1$ vector of constants with $n = m + k$, $\boldsymbol{B}_l$ for $l = 1, ..., p$ are $n \times n$ coefficient matrices and $\boldsymbol{\nu}_t$ is a $n \times 1$ vector of normally distributed reduced form errors with zero mean and variance-covariance matrix $\boldsymbol{\Sigma}$. Each reduced form error can be expressed as a linear combination of structural shocks $\boldsymbol{u}_t$, i.e. $\boldsymbol{\nu}_t = \boldsymbol{G}\boldsymbol{u}_t$, which implies that $\boldsymbol{\Sigma} = \boldsymbol{G}\boldsymbol{G}'$, where $\boldsymbol{u}_t$ is an $n \times 1$ vector of normally distributed structural shocks with zero mean and identity variance-covariance matrix. Moreover, Equation (1) can be rewritten as

$$y_{j,t} = \boldsymbol{\beta}_j' \boldsymbol{Z}_t + \epsilon_{j,t} \tag{3}$$

where $\boldsymbol{\beta}_j = [\boldsymbol{\beta}_j^x \ \boldsymbol{\beta}_j^\kappa]'$ is an $n \times 1$ vector of coefficients.

The identification of structural shocks with a clear economic interpretation from reduced form residuals that do not have any economy economic interpretation is a key step in VAR models. Here we identify shocks by imposing some restrictions on the sign of impulse responses on impact (as recommended by Canova and Paustian (2011), following the procedure described in Rubio-Ramirez, Waggoner and Zha (2010).[5] The aim of this exercise is to extract the exogenous component of the immigration series, while controlling for its endogenous response to, for example, changes in economic activity. Our identification procedure thus mainly focuses on disentangling business cycle from immigration shocks. To identify an immigration shock, we restrict on impact the sign of the impact impulse response of employment and immigration. For the business cycle shock, we restrict employment, immigration, real wages and the participation rate. To achieve identification, a magnitude restriction plays a critical role in our model, in a re-interpretation of the max share approach originally proposed by Uhlig (2004). Intuitively, we impose that employment (a cyclical variables in Norway) should be mainly (but not exclusively) driven by business cycle shocks while the immigration variable should be driven mainly (but not

[5]Further details on the estimation of VAR with sign restrictions, including the specification of the priors, can be found in Furlanetto and Robstad (2019).

Table 1: *Identification Restrictions*

	Business Cycle	Immigration
Employment	+	+
Immigration	+	+
Labor Force Participation	+	NA
Real Wages	+	NA
Employment per Immigrant	+	-

The table describes the sign restrictions on the impact impulse response function used for each variable or ratio (in rows) to shocks (in columns). NA indicates that the response of the variable is left unrestricted.

exclusively) by immigration shocks. The strategy is implemented by imposing a simple sign restriction over the ratio of employment over immigration. While both shocks are expected to move both variables in the same direction, it is reasonable to assume that a business cycle shock will move employment more than immigration while an immigration shock will move more immigration than total employment. Table 1 summarizes our restrictions.

Given that the adjacent income percentile possibly share similar dynamics, we extend the framework of Bernanke, Boivin and Eliasz (2005) by connecting the coefficients of the adjacent percentiles.[6] This approach allows for further dimension reduction. We use the following structure on the coefficients

$$\boldsymbol{\beta}_j = \boldsymbol{\alpha} + \sum_{i=1}^{q} \boldsymbol{\Theta}_i \boldsymbol{\beta}_{j-i} + \boldsymbol{\zeta}_j \tag{4}$$

where $\boldsymbol{\alpha}$ is an $n \times 1$ vector of constants, $\boldsymbol{\Theta}_i$ for $i = 1, ..., q$ are $n \times n$ diagonal matrices with autoregressive coefficients on their diagonals, $\boldsymbol{\zeta}_j$ is normally distributed error with zero mean and diagonal variance-covariance matrix $\boldsymbol{\Omega}$. In the estimation exercise we truncate the autoregressive coefficients to the stationary region.

We estimate the model using standard Bayesian methods. To do so, we cast the model in the state space form, where Equation (3) serves as the observation equation and Equation

[6]In contrast to our approach, Bernanke, Boivin and Eliasz (2005) extract their factors from various macroeconomic and financial variables, which do not exhibit these kind of "spatial" correlations.

(2) as the state equation. All priors used are relatively diffuse. We use Gibbs sampling to simulate from the posterior of the unobserved states and parameters (see, e.g., Kim and Nelson, 1999 for a textbook discussion).

The Gibbs sampler consists of three blocks. In the first block, we condition on the unobserved states $\boldsymbol{\kappa} = \{\boldsymbol{\kappa}_t\}_{t=1}^{T}$ and $\boldsymbol{\beta} = \{\boldsymbol{\beta}_j\}_{j=1}^{99}$ and draw the model's parameters. More precisely, we draw the VAR parameters $\boldsymbol{c}$, $\boldsymbol{B}_l$ and $\boldsymbol{\Sigma}$ from a normal-inverse-Wishart distribution, the AR parameters in $\boldsymbol{\Theta}_i$ and the variances in $\boldsymbol{\Omega}$ from a Normal-inverse-Gamma distribution and the variances $\sigma_1^2, ..., \sigma_{99}^2$ from an inverse Gamma distribution. In the second block, we condition on the parameters of the model and the unobserved states $\boldsymbol{\beta}$ and draw $\boldsymbol{\kappa}$ using the forward filtering-backward sampling approach described in Carter and Kohn (1994) and Frühwirth-Schnatter (1994). In the third block, we condition on all the parameters and the states $\boldsymbol{\kappa}$ and draw $\boldsymbol{\beta}$ using again the forward filtering-backward sampling algorithm. We simulate 100k draws from the Gibbs sampler, discarding the first 80k draws as burn-in and saving every 10th draw, which results into 2000 draws that we use for inference. We checked convergence of the Gibbs sampler by visual inspection, i.e. by consulting recursive mean plots and by using the statistical testing procedure suggested by Geweke (1992). In the empirical exercise below we specify $p = 2$ for the lag length of VAR and $q = 1$ for the lag length of AR-process in (4), $k = 1$ factors and $m = 4$ macroeconomic variables.

4 Empirical Results

In this Section, we present the results obtained from the estimation exercise. In a first step, we show impulse responses and variance decompositions for the aggregate variables. In a second step, we focus on the distributive effects of immigration shocks and plot impulse responses for the percentiles of the earnings distribution and for a summary measure of inequality.

Figure 4: IRF business cycle shock

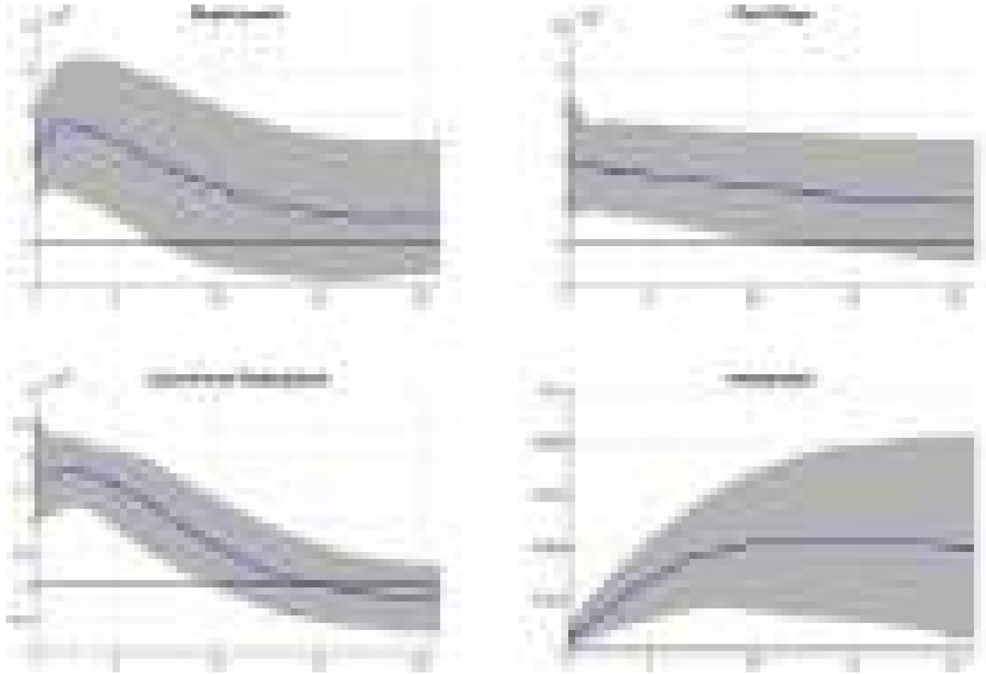

Notes: The dash-dotted blue line depicts the posterior median and the grey shaded area represents the 68% error bands for the impulse responses to an one standard deviation business cycle shock. The estimation sample spans the period 1997Q1–2019Q4.

We begin our analysis with the model using aggregate data on immigration (thus mixing all reasons for immigration together). In Figures 4 and 5 we present impulse responses to the business cycle shock and to the immigration shock. Not surprisingly, the business cycle shock has persistent effects on the cyclical variables included in the system i.e. employment, real wages and the labor force participation rate. More interestingly, we remark that the stock of immigrants gradually builds up and peaks after around 10 quarters. Thus, immigration responds endogenously to the cycle. The immigration shock has a large impact effect on the stock of immigrants which peaks after 7 quarters and a non-negligible effect also on aggregate employment.

Figure 5: IRF immigration shock

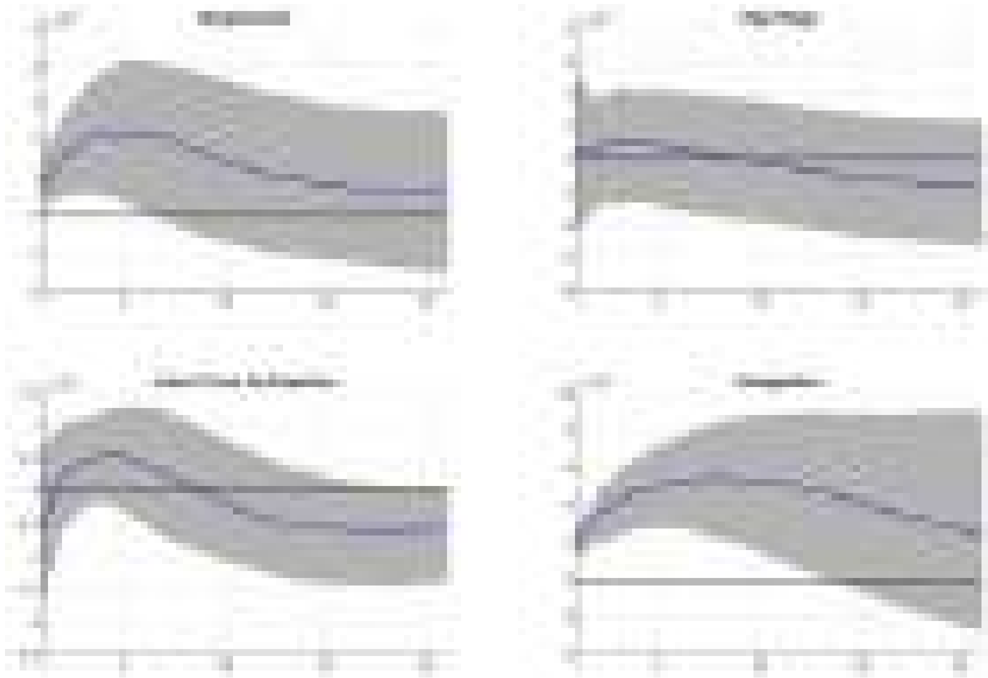

Notes: The dash-dotted blue line depicts the posterior median and the grey shaded area represents the 68% error bands for the impulse responses to an one standard deviation immigration shock. The estimation sample spans the period 1997Q1–2019Q4.

Interestingly, we detect some effects also on the labor force participation rate which exhibits a tiny increase after one year. Presumably, such a tiny effect is an average across different migration shocks: intuitively, we expect the participation rate to increase in response to job-related immigration shocks while it is conceivable to imagine a decline (at least on impact) in response to a refugee shock. Aggregate real wages are unrestricted in the system and their response in imprecisely estimated. The point-wise median across all draws satisfying the sign restrictions lies around zero, in keeping with the conventional wisdom that immigration shocks have limited effects on average wages in the economy. The variance decomposition presented in Figure 6 presents the same information into a different format. The blue areas show the share of variance explained by the business cycle shock for each variable at various horizons. The green areas summarize the importance of immigration shocks. In keeping with our identification assumptions, immigration is mainly driven by immigration shocks while employment is mainly driven by business cycle shocks. However, it is interesting to notice how the impact of immigration on employment is not negligible at all, unlike for wages. Thus, immigration shocks contribute to business cycle fluctuations in Norway. In contrast, immigration is a largely exogenous phenomenon according to the idea that immigration is largely driven by changes in legislation and economic conditions abroad more than from economic conditions in Norway.

Figure 6: Variance decomposition baseline model

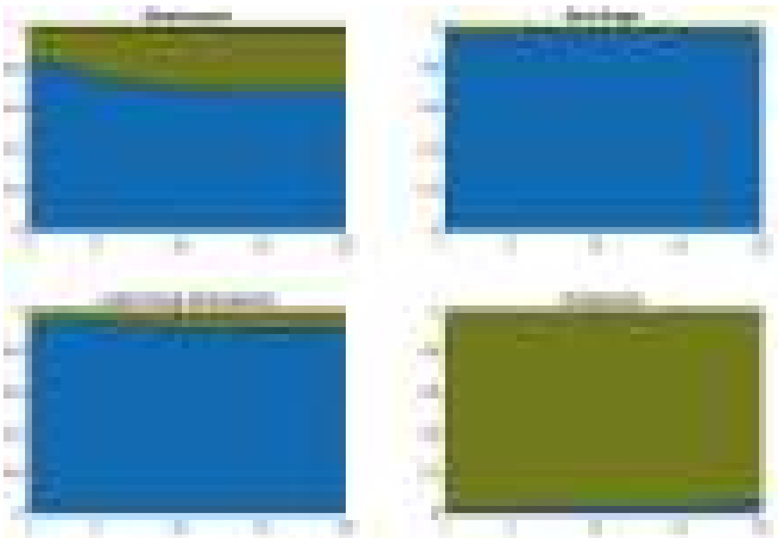

Notes: Median forecast error variance decomposition. The blue area represent the fraction of forecast error variance explained by the business cycle shock. The green area depicts the fraction explained by the immigration shock.

We now turn to the distributional effects of immigration in our system. In Figure 7 we summarize the effects of immigration shocks across the percentiles of the earnings distributions. The vertical axis refers to the magnitude of the effects, the left horizontal axis indicates the horizon of the impulse response functions while the right horizontal axis refers to the percentiles with 1 indicating the percentile with lowest earnings and 99 indicating the percentile with the highest earnings. The response of the percentiles around the median is extremely small, slightly on the negative side, insignificant both statistically and economically.

Figure 7: IRF immigration shock bird's eye view

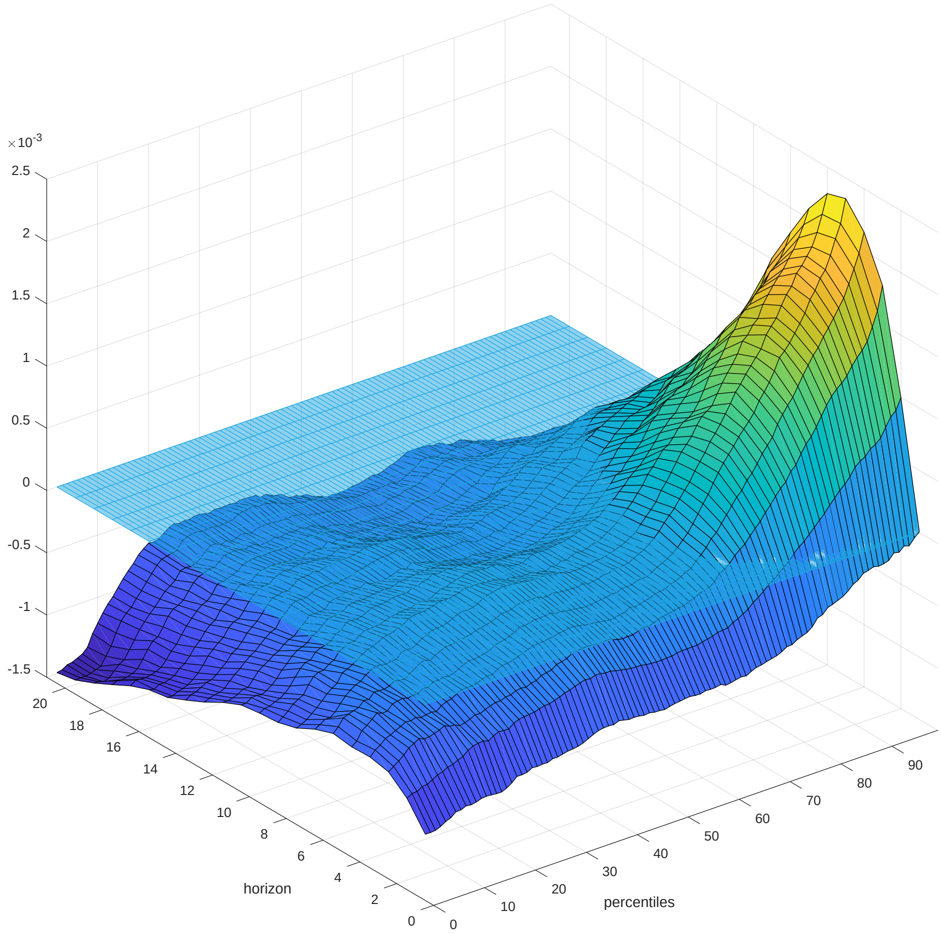

Notes: 3-D representation of the posterior median impulse responses to a one standard deviation immigration shock.

This result on earnings confirms the result on average wages discussed earlier in keeping with the previous literature. Notably, however, the effects are larger towards the tail of the distribution. The response of the lowest percentiles is more negative, especially at longer horizons, while the response of the highest percentiles in the distribution is rather large and positive. Thus according to our model, the effects of immigration shocks on earnings are rather small but heterogeneous and lead to an increase in labor income inequality, mainly driven by the right tail of the distribution. At this stage, it is important whether such an increase in inequality is significant from a statistical points of view. In order to tackle this question we focus our attention on a summary measure of inequality

Figure 8: Percentile 90 minus percentile 10 (P90-P10)

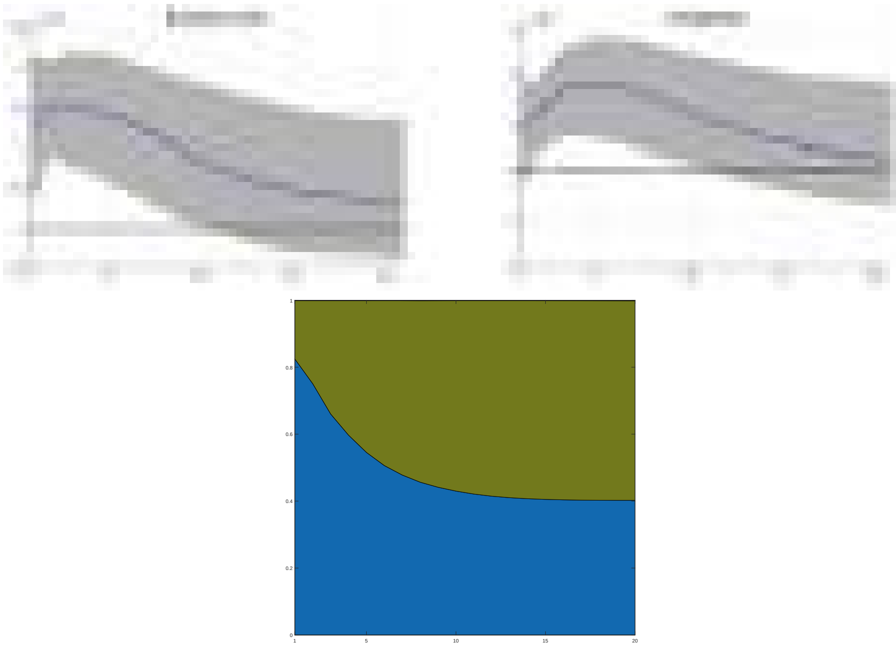

Notes: The upper left panel shows impulse responses of P90-P10 to a one standard deviation shock in the busines cycle shock and the upper right panel shows the responses of P90-10 to the immigration shock. The lower panel reports median forecast error variance decomposition. The blue area represent the fraction of forecast error variance explained by the business cycle shock. The green area depicts the fraction explained by the immigration shock.

(P90 - P10) commonly used in the literature and we plot its impulse response function with 68 per cent error bands obtained from the 2000 draws of the impulse responses for P90 and P10.

In Figure 8 we show that the increase in inequality is rather large, statistically significant according to 68 per cent error bands and peaking after 5 quarters. Two comments are in order. First, the increase in inequality generated by immigration is comparable to the one generated by the business cycle shock, a result confirmed also by the variance decomposition which shows how immigration explain 60 per cent of the variation in P90 - P10 at long horizons (while business cycle shocks are more important in the short run).[7] Second, this measure of inequality is pro-cyclicality in response to both shocks. This means that inequality increases in booms and decreases in recessions unlike in the

[7] Put simply, an immigration shock that increases the immigrant population by about 6 percent on impact would lead to an increase in the earnings gap between the 90th and 10th percentile by about 2 percent at the peak. In 2019 this would correspond to about 50 000 immigrants and a increase in the wage difference of about 24 000 NOK (2 400 dollars) in annual wage terms.

US where earning inequality increases in recessions driven by the negative effects on the left tail of the distribution. Such a result has potentially important macroeconomic implications since the amplification in the propagation of several macroeconomic shocks in New Keynesian models with simple forms of heterogeneity relies on the presence of countercyclical inequality (cf. Bilbiie (2020)).

Finally, we compare our baseline model with an alternative model using the earnings distribution for native workers only. Figure 9 compares the impulse response of P90 - P10 in the two models. We see that the effect is confirmed but substantially attenuated when we focus only on native workers. This hints to the fact that non-native workers (i.e. previous immigrants) are disproportionately affected by immigration shocks and contribute substantially to the increase inequality. Therefore, we conform this classic result of the immigration literature that has been found recently also by Maffei-Faccioli and Vella (2021) using German macro data.[8]

[8] In this first draft of the paper we dissect the earnings distribution only in the native/non native dimension. Note, however, that it is conceivable to run exercises also along the gender, age, education, sector of specialization dimensions. One specific exercise that we have in mind is to compare the response of the earnings distributions for employees in the private sector and employees in the public sector (which is large in Norway in comparison with international standards). Displacement effects in earnings seems more natural to arise in the private sector where the share of immigrants employed is substantially higher.

Figure 9: P90 minus P10 for natives and total population

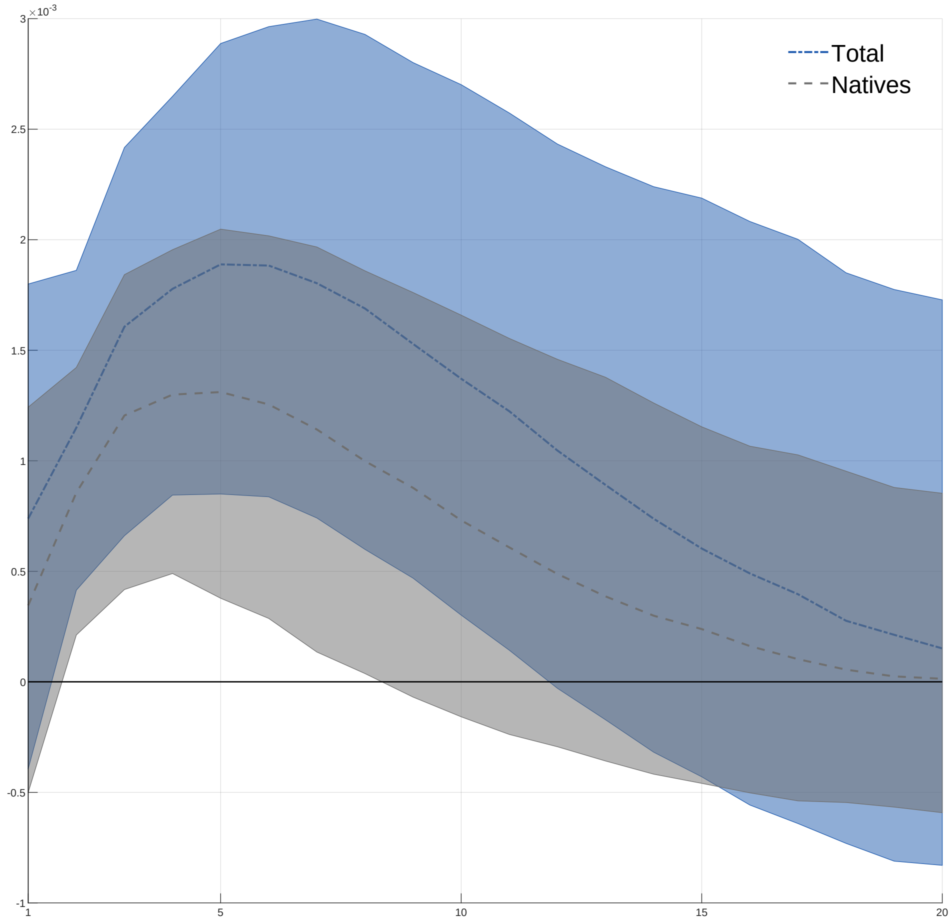

Notes: The dash-dotted blue line depicts the posterior median impulse response of P90-P10 to a one standard deviation shock in the immigration shock for the total population and the red dashed line is the response of the natives only. The shaded areas represent the 68% error bands.

5 Extensions: reason of immigration, education and country of origin

In this Section, we dissect the immigration series along several dimensions and we consider the effects of more granular immigration shocks. A nice property of our admittedly simple identification scheme is that it can be maintained in each extension. In fact, we impose only that immigration shocks must have a large positive impact effect on the immigration series used in each experiment and a positive impact effect on aggregate employment.

The advantage of using the immigration series derived from the microdata (and plotted in Figure 10) is that we can slice it in several dimensions. Here we will consider various immigration shocks zooming in on the reason for immigration, the level of education and the country of origin of immigrants. Note, however, that we could have considered also other dimensions like gender (by comparing a female immigration shock against a male immigration shock), age (by comparing an inflow of young workers against an inflow of older workers) or sector of specialization (by comparing, for example, an inflow in the construction sector against an inflow in the hospitality sector). We plan to perform some of these experiments in the future.[9]

Figure 10: Different sources of immigration

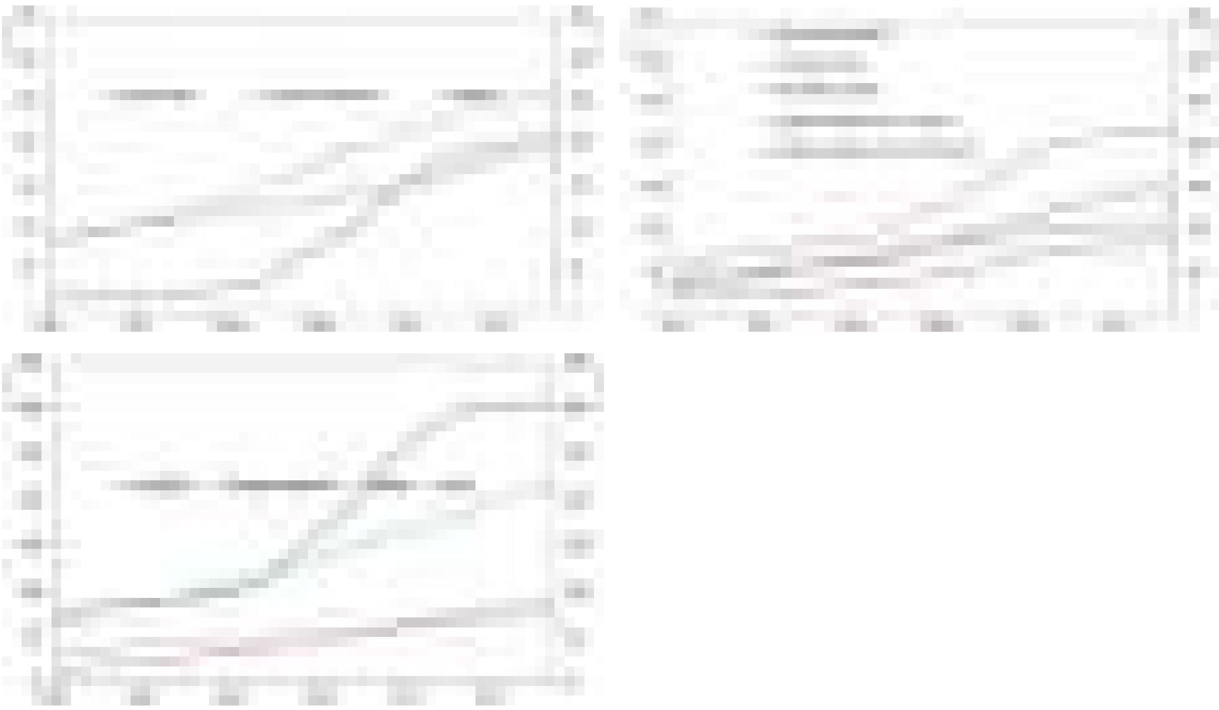

In a first experiment, we consider three immigration shocks by considering the reason for immigration. Statistics Norway records five different reasons for immigration by

[9]In addition, the series could be sliced further and across several dimensions at the same time. In principle, we could consider the effects of, let say, a job related immigration flow of high skilled immigrants from Germany against a job related flow of high school dropouts from Poland working in the construction sector. The disadvantage of these experiments, which are per se very intriguing, is that the macroeconomic impact of the impulse diminishes with the granularity of the experiment. A more microeconomic approach is probably more suitable in these cases.

distinguishing job related immigration, refugees, family reunifications, students and a residual category. We focus our attention on the first three groups that account for the bulk of migration flows in recent years as shown in Figure 10. We focus our attention on the response of P90 - P10 in response to the three immigration shocks and we plot the peak value of each response (which happens around horizon 8 in all cases) with 68 per cent credibility bands in Figure 11. We see that job related immigration has by far the largest effect on inequality, thus driving the result obtained using the aggregate immigration series. Family reunifications have an intermediate effect on inequality with the credibility bands including the value of zero at the margin while the effect induced by refugees immigration is lower and insignificant, both statistically and economically. We believe this result confirms the macroeconomic importance of job-related immigration which was the focus of our previous paper in a model with tighter identification assumptions (cf. Furlanetto and Robstad (2019).

Figure 11: P90 minus P10: reason for immigration

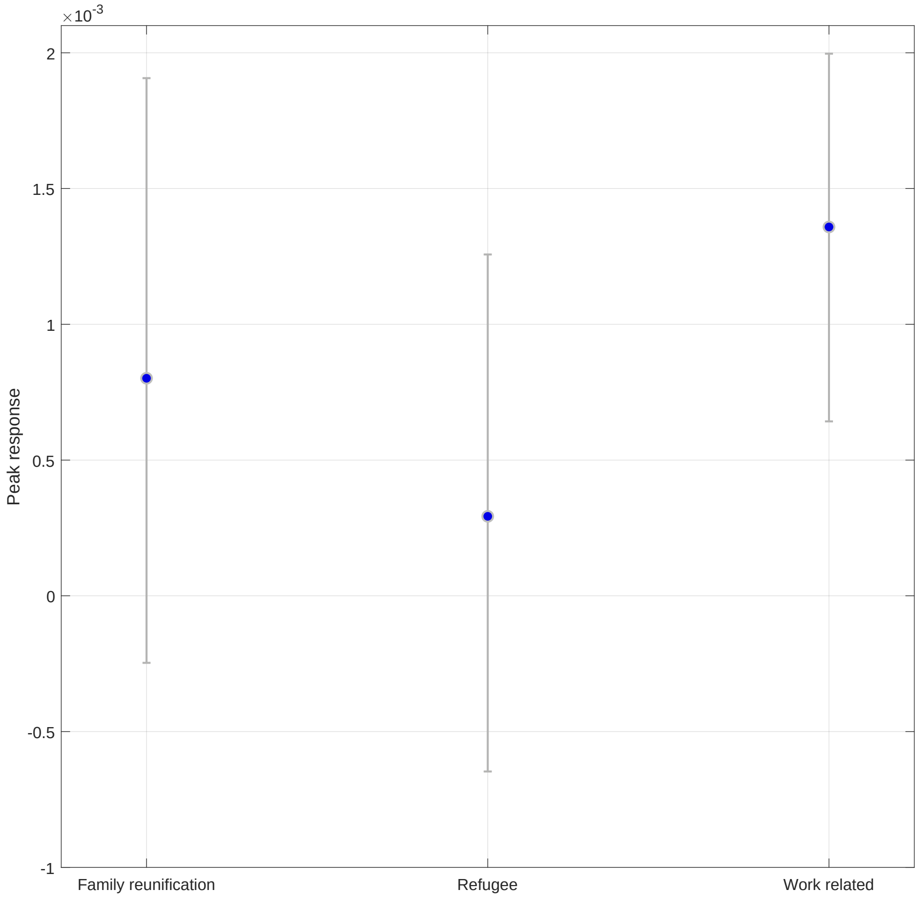

In a second experiment, we consider the education level of immigrants which are classified into five groups. Group 1 includes immigrants with no formal education completed, group 2 refers to basic school completed, group 3 to completion of secondary school, group 4 to up to four years of university education and group 5 to more then 4 years of university education. As shown in Figure 10, immigrants in group 3 account for the largest group. In Figure 12 we show that inequality increases in response to every kind of immigration shock, independently from the education level of immigrants. The magnitude of the effects, however, depends on the level of education. Group 3 has by far the largest effect on inequality but non negligible effects are induced also by groups 1, 4 and 5. The lowest effects are driven by Group 2.

Figure 12: P90 minus P10: Education

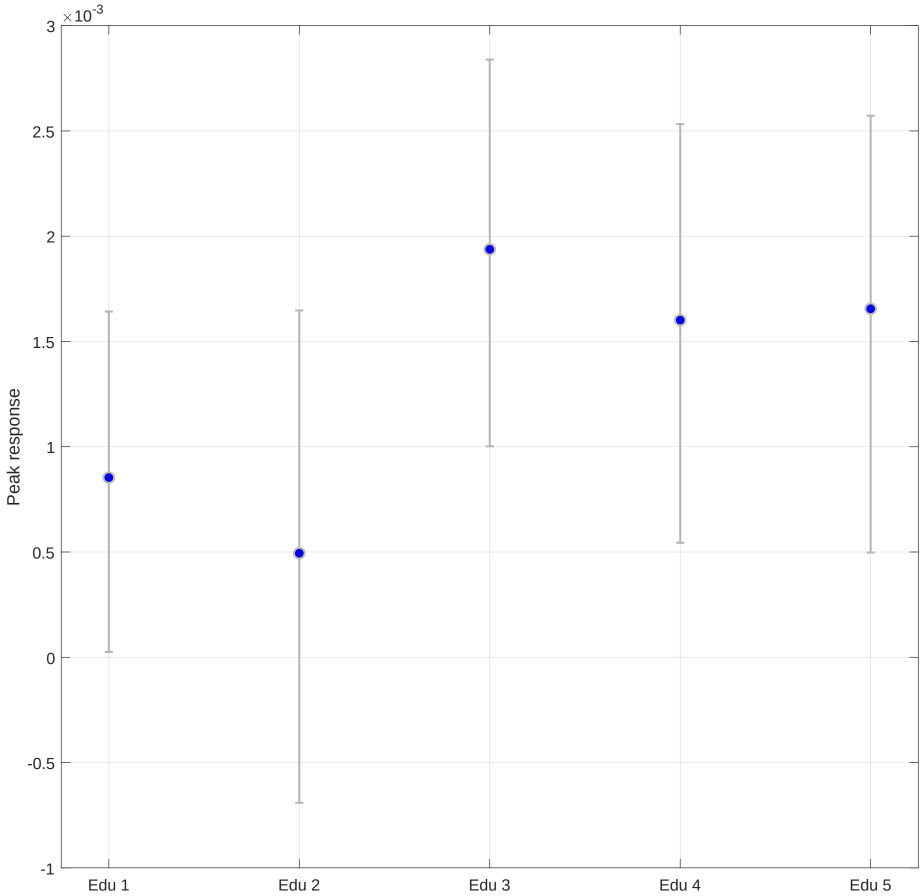

In a third experiment, we consider the country of origin of immigrants and we focus on the four largest country groups: immigrants from the European Union (EU) and European Economic Area (EEA), remaining European countries, Asia and Africa. Once again, the effects on earnings inequality are rather heterogeneous with the largest impact induced by immigrants from the EU/EEA group which is dominated by Poles, Swedes and Lithuanians. Most of these workers are job-related immigrants with an intermediate level of education. Therefore, all in all, the three more granular exercises seem to point in the same direction. The kind of immigration that generates an increase in earnings inequality is essentially the immigration induced by the EU enlargements to Eastern European countries that took place over the last 20 years.

Figure 13: P90 minus P10: Country of origin

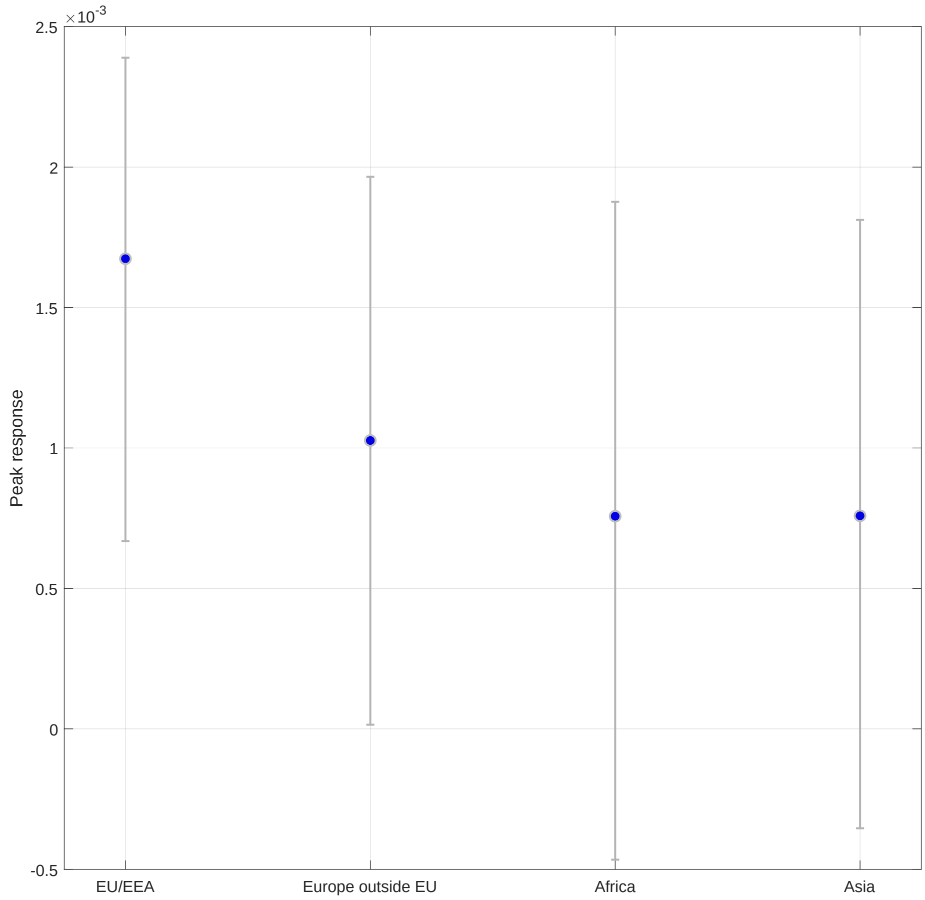

6 Conclusion

In this paper, we have investigated the impact of immigration on labor earnings inequality. Our main result is that immigration shocks increase earnings inequality with the effect mainly driven by workers with highest earnings who benefits disproportionately from the shock. A possible interpretation for this result refers to complementarities between immigrants and rich natives. The increase in inequality is induced mainly by job related immigration, by immigrants with intermediate to high education level and by immigrants from EU/EEA. All these characteristics seem to be consistent with the large migration flows induced by EU enlargements in recent years.

Figure 14: IRF job related immigration shock bird's eye view

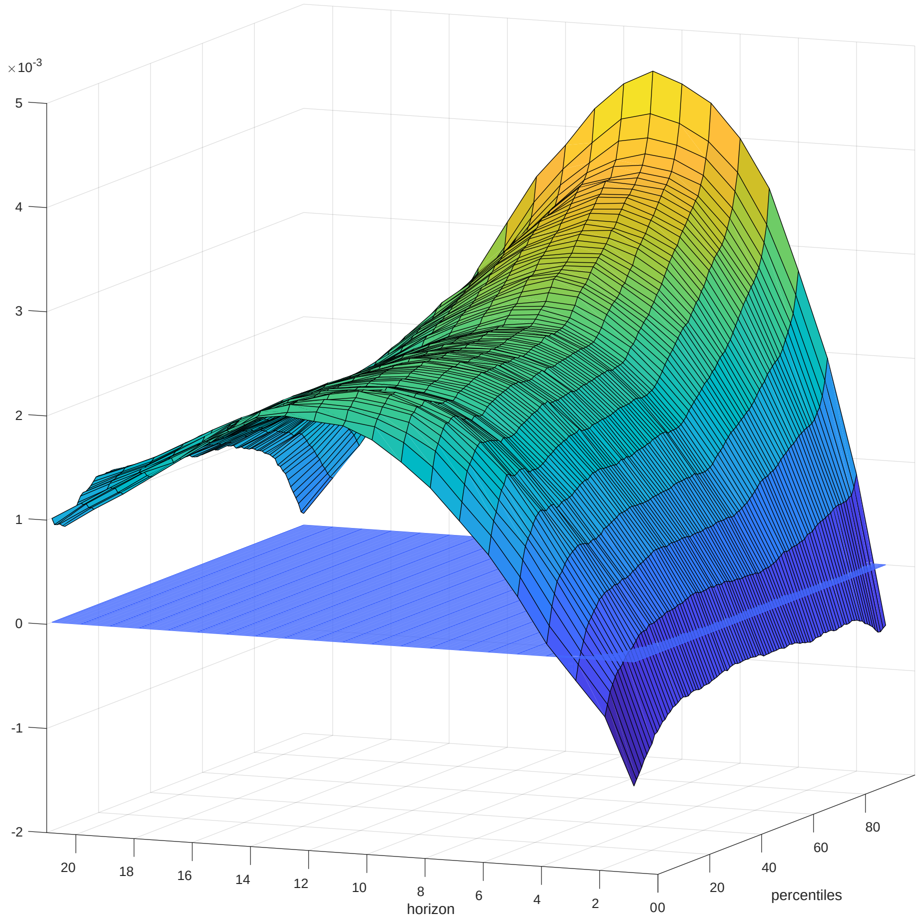

Notes: 3-D representation of the posterior median impulse responses to a one standard deviation job related immigration shock.

It is of paramount importance to stress that our analysis is purely positive and that no normative implications should be extracted from our work. The increase in inequality induced by immigration is not necessarily worrisome for the Norwegian society. In fact, the rise in inequality is substantially lower when we focus only on native workers in keeping with the idea that immigrants compete first and foremost with previous immigrants. In addition, the impact of immigration shocks on the *level* of earnings seems to be rather benign across the distribution with most workers being barely affected by the shock and with rich workers being positively affected. In this paper, we have not stressed the impact of the shock on the level of earnings because it is imprecisely estimated given that we impose very few identification assumptions (for the case of job related immigration the effects seem to be more expansionary at the bottom of the distribution, as shown in Figure 14, but once again the uncertainty around these effects is too large to make strong statements on the *level* effect). In the coming months, we plan to refine the identification scheme by imposing perhaps some additional assumptions in order to derive more precise results on the level effect. Notably, however, our relatively agnostic identification scheme is sufficient to derive reasonably robust results on the inequality dimension.

References

Basso, Gaetano, Giovanni Peri, and Ahmed S Rahman. 2020. "Computerization and immigration: Theory and evidence from the United States." *Canadian Journal of Economics*, 53(4): 1457–1494.

Bernanke, Ben S, Jean Boivin, and Piotr Eliasz. 2005. "Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach." *The Quarterly Journal of Economics*, 120(1): 387–422.

Bilbiie, Florin O. 2020. "The new keynesian cross." *Journal of Monetary Economics*, 114: 90–108.

Bratsberg, Bernt, and Oddbjørn Raaum. 2012. "Immigration and wages: Evidence from construction." *The Economic Journal*, 122(565): 1177–1205.

Bratsberg, Bernt, Oddbjørn Raaum, and Knut Røed. 2014. "Immigrants, labour market performance and social insurance." *The Economic Journal*, 124(580): F644–F683.

Canova, Fabio, and Matthias Paustian. 2011. "Business cycle measurement with some theory." *Journal of Monetary Economics*, 58(4): 345–361.

Card, David. 2009. "Immigration and inequality." *American Economic Review*, 99(2): 1–21.

Carter, C., and R. Kohn. 1994. "On Gibbs sampling for state space models." *Biometrika*, 81: 541–553.

Chang, Minsu, and Frank Schorfheide. 2022. "On the Effects of Monetary Policy Shocks on Earnings and Consumption Heterogeneity."

Chang, Yoosoon, Fabio Gómez-Rodríguez, and Gee Hee Hong. 2022. "The Effects of Economic Shocks on Heterogeneous Inflation Expectations."

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary policy and inequality." *Journal of Monetary Economics*, 88: 70–89.

De Giorgi, Giacomo, and Luca Gambetti. 2017. "Business cycle fluctuations and the distribution of consumption." *Review of Economic Dynamics*, 23: 19–41.

d'Albis, Hippolyte, Ekrame Boubtane, and Dramane Coulibaly. 2019. "Immigration and public finances in OECD countries." *Journal of Economic Dynamics and Control*, 99: 116–151.

d'Albis, Hippolyte, Ekrame Boubtane, and Dramane Coulibaly. 2021. "Demographic changes and the labor income share." *European Economic Review*, 131: 103614.

Frühwirth-Schnatter, Sylvia. 1994. "Data augmentation and dynamic linear models." *Journal of Time Series Analysis*, 15: 183–202.

Furlanetto, Francesco, and Ørjan Robstad. 2019. "Immigration and the macroeconomy: Some new empirical evidence." *Review of Economic Dynamics*, 34: 1–19.

Geweke, John. 1992. "Evaluating the accuracy of sampling-based approaches to calculating posterior moments." In *Bayesian Statistics 4.* , ed. M Bernado, J.O. Berger, A.P. Dawid and A.F.M. Smith, 169–193. Oxford, UK:Clarendon Press.

Hauser, Daniela, and Martin Seneca. 2022. "Labor mobility in a monetary union." *Journal of International Economics*, 137: 103600.

Hoen, Maria F, Simen Markussen, and Knut Røed. 2022. "Immigration and economic mobility." *Journal of Population Economics*, 35(4): 1589–1630.

Jentoft, Susie. 2014. "Imputation of missing data among immigrants in the Register of the Population's Level of Education (BU)."

Kerr, Sari Pekkala, and William R Kerr. 2011. "Economic impacts of immigration: A survey."

Kiguchi, Takehiro, and Andrew Mountford. 2019. "Immigration and unemployment: A macroeconomic approach." *Macroeconomic Dynamics*, 23(4): 1313–1339.

Kim, Chang-Jin, and Charles R. Nelson. 1999. *State-Space Models with Regime Switching.* Cambridge MA and London:The MIT Press.

Lewis, Ethan. 2011. "Immigration, skill mix, and capital skill complementarity." *The Quarterly Journal of Economics*, 126(2): 1029–1069.

Maffei-Faccioli, Nicolò, and Eugenia Vella. 2021. "Does immigration grow the pie? Asymmetric evidence from Germany." *European Economic Review*, 138: 103846.

Olovsson, Conny, Karl Walentin, and Andreas Westermark. 2021. "Dynamic macroeconomic implications of immigration." Sveriges Riksbank Working Paper Series.

Ottaviano, Gianmarco IP, and Giovanni Peri. 2012. "Rethinking the effect of immigration on wages." *Journal of the European Economic Association*, 10(1): 152–197.

Rubio-Ramirez, Juan F., Daniel F. Waggoner, and Tao Zha. 2010. "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference." *Review of Economic Studies*, 77(2): 665–696.

Schiman, Stefan. 2021. "Labor supply shocks and the Beveridge Curve—Empirical evidence from EU enlargement." *Review of Economic Dynamics*, 40: 108–127.

Smith, Christie, and Christoph Thoenissen. 2019. "Skilled migration and business cycle dynamics." *Journal of Economic Dynamics and Control*, 109: 103781.

Uhlig, Harald. 2004. "What moves GNP?" Econometric Society.