

Ruisi, Germano

Working Paper

An assessment of the macroeconomic implications of foreign and domestic labour supply shocks in Malta

CBM Working Papers, No. WP/06/2020

Provided in Cooperation with:

Central Bank of Malta, Valletta

Suggested Citation: Ruisi, Germano (2020) : An assessment of the macroeconomic implications of foreign and domestic labour supply shocks in Malta, CBM Working Papers, No. WP/06/2020, Central Bank of Malta, Valletta

This Version is available at:

<https://hdl.handle.net/10419/237306>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



BANK ĊENTRALI TA' MALTA
EUROSISTEMA
CENTRAL BANK OF MALTA

An Assessment of the Macroeconomic Implications of Foreign and Domestic Labour Supply Shocks in Malta

Germano Ruisi¹

WP/06/2020

¹Senior Research Economist, Research Department, Central Bank of Malta.
Email: ruisig@centralbankmalta.org

I would like to thank Alexander Demarco, Aaron Grech, Brian Micallef, Noel Rapa, William Gatt Fenech and Črt Lenarčič (Bank of Slovenia) for their useful comments and suggestions. I am also thankful to Abigail Marie Rapa, Reuben Ellul and John Farrugia for assistance with data collection. Any remaining errors are my own. The views expressed in this paper are the author's own only and do not reflect those of the Central Bank of Malta.

Abstract

Over the recent years Malta has experienced a remarkable increase in its labour force due to a large influx of immigrants and an unprecedented increase in the domestic participation. Driven by the observation of such a phenomenon, this paper aims at assessing the impact of foreign and domestic labour supply shocks on the Maltese economy by estimating a number of structural vector autoregressions (VARs) identified through sign restrictions. The VARs are estimated by using data over the 2004Q1-2019Q2 period and the results point toward a relevant impact of the identified shocks on domestic production, wages and unemployment as well as on a number of other key variables, e.g., government revenue and expenditure, rents and measures of productivity.

JEL Classification: C11, C32, E32

Keywords: Labour Supply Shocks, Structural VARs, Sign Restrictions

Executive Summary

In the aftermath of the 2009 global financial crisis, the Maltese labour market exhibited remarkable resilience, even in the face of adverse external shocks, with a decline in the unemployment rate to historical lows even as the labour participation rates rose sharply. Two factors were especially important in increasing the labour supply: rising female participation rates and an influx of foreign workers. The female participation rate for those aged between 16 to 64 years of age has quickly converged to the EU average, increasing from 36% in 2005 to 63.75% in 2018. At the same time, the share of foreign workers rose from 0.25% of the total working age population aged 16 to 64 at the time of EU membership in 2004 to 16.39% in 2018. In both cases, these two factors had a significant impact on the economy's potential output growth (Grech (2017); Micallef (2018)).

This paper focuses on the macroeconomic implications of labour supply shocks, both foreign and domestic, on the Maltese economy using a Bayesian vector autoregression model (BVAR). The identification strategy, which relies on sign restrictions, is used to disentangle five structural shocks: two labour supply shocks, which are the focus of this paper, a wage bargaining shock, a business cycle shock and another catch-all shock to close the system. The restrictions are mainly derived from a set of theoretical models and closely follow the approach used in Furlanetto & Robstad (2019). The model is estimated over the period 2004Q1-2019Q2 on five macroeconomic variables: real GDP, real wages, the unemployment rate, the participation rate and the share of migrants relative to the total working age population.

A foreign labour supply shock normalised to increase in the immigration share by 1% on impact has a positive and persistent effect on real economic activity that lasts for around four years. The response of the immigration share is humped-shaped, reaching a peak after around a year. The increase in the foreign component of the labour supply raises the participation rate, lowers the unemployment rate and generates downward pressure on real wages. The impact on the latter two variables is mostly significant over the medium-term. The impulse responses from an increase in the domestic labour supply are broadly similar, both quantitatively and qualitatively, with two exceptions. First, it lowers the share of migrants and second, it has a positive, though short-lived effect on impact on the unemployment rate. The latter is, however, not significant and

the impact on unemployment in the medium term is negative.

According to the forecast error variance decomposition, the two identified labour supply shocks contribute in a remarkable way to the unexpected fluctuations of the domestic variables. For real GDP, the two labour supply shocks explain roughly 35% of fluctuations on impact and around 60% after ten years. The business cycle shock explains around 20% of GDP on impact and 35% in the long run. The two labour supply shocks and the business cycle shocks also have a predominant effect on real wages and the unemployment rate.

A sensitivity analysis is used to assess the robustness of the baseline results to a change in the identification strategy. The focus is on the impact on real wages by relaxing the assumption that wages have to fall on impact as soon as the labour supply shocks hit the economy (more precisely the assumption is that wages do not move on impact). This is motivated by sticky wages and by the fact that, especially for migration, the impact crucially depends on the skill composition of migrants and natives. Furthermore, some studies that rely on micro data from other countries find that migrants and natives are likely to specialise and perform differentiated tasks, thus complementing each other even within the same skill category (Peri & Sparber (2009); Peri (2016)). The results suggest that real wages might also increase in the short-run following an immigration shock, in line with the evidence from the micro literature, though the impact over the medium-term is quantitatively and qualitatively similar to the baseline model. A similar finding is obtained with the domestic labour supply shock, though the short-run effect on wages is not different from zero.

The paper also assesses the relevance of the two labour supply shocks on nine macroeconomic variables that broadly capture the effects on public finances, consumer price inflation, the housing market and productivity. Each exercise is conducted by substituting the unemployment rate in the baseline model with one variable of interest at a time, which is designed to ensure that the model remains fully identified. The results are broadly in line with the evidence from the literature (MAC (2018)) though in some cases there is a considerable degree of uncertainty surrounding the estimates. In terms of the immigration shock, the impact on public finances is positive as the effects on government revenues outweigh those on expenditure. No clear cut evidence is obtained from consumer price inflation. Foreign labour supply shocks have a positive effect on the housing market, though the impact on rents is more pronounced and significant compared to house prices. Finally, immigration shocks exerted a positive impact on

labour productivity. The latter effect was entirely driven by an increase in total factor productivity, whereas the impact on the capital intensity is negative. With few exceptions, the results for domestic labour supply shocks are qualitatively similar. The increase in productivity, in line with Micallef & Ellul (2017), explains the importance of the increase in labour supply in addressing crucial skills and labour shortages, which were key to raise the country's supply-side potential.

Contents

1	Introduction	1
2	Methodology	3
2.1	Model	3
2.2	Identification Strategy	4
3	Results	6
3.1	Baseline Results	7
3.2	Further Experiments on Selected Variables of Interest	13
3.3	Sensitivity	19
4	Conclusions	21
A	Appendix	25
A.1	Data Sources and Transformations	25
A.2	Estimation Details	26
A.3	Lag Selection	29
A.4	Convergence Diagnostics	30
A.5	Responses to Non Labour Supply Shocks	31

1 Introduction

Over the recent years Malta has experienced a remarkable increase in its labour force driven by both a large influx of foreign workers and a exceptional increase in the participation of Maltese people. The increase in the Maltese labour force has become more and more pronounced since 2012 as shown in figure 1 where the dynamics of the Maltese labour market are depicted.

The top left panel of figure 1 shows the evolution of the number of foreign workers compared with the Maltese total working age population of those aged between 16 and 64 years. The amount of foreign workers was negligible until 2004. After this date, with the Maltese accession to the European Union and the higher mobility of people that the latter brought about, the number of foreign workers in the economy started to rise and experienced a noticeable acceleration after 2012. In addition, a number of policies aiming at allowing third country nationals to join the Maltese labour force have been recently implemented by the government. The top right panel of figure 1 shows how the share of foreign workers rose from 0.25% of the total working age population aged 16 to 64 at the time of EU membership in 2004 to a value of 16.39% in 2018. This accounted to a stock of foreign workers of nearly 53,000 out of a total working age population of around 325,000. The large influx of immigrants that the Maltese economy has experienced, however, is not an isolated phenomenon as similar dynamics have been seen in most advanced economies. For example, the share of immigrants in 2016 was equal to 13.3% for the United Kingdom, 11.8% for France, 13.3% for Germany and 14.9% for Norway. As opposed to the mentioned economies, that have long been recipients of migrants, what is striking for the Maltese case is the rapidity with which such a phenomenon has taken place.

The two bottom panels of figure 1 aim at displaying how the Maltese workforce and the participation rate have increased over time. The bottom left panel demonstrates how the latter variables have both been characterised by an upward trending behaviour since 2004 with an acceleration since 2012. The overall participation rate was around 59% in 2004 but then it went above 74% in 2018. The bottom right panel suggests how this was particularly driven by an acceleration in the female participation. The latter, for those aged between 16 to 64 years of age, has quickly converged to the EU average, increasing from 36% in 2005 to 63.75% in 2018. As opposed to that, the male participation rate has shown much more contained changes over time with values

ranging from 77% to 84.5% respectively in 2009 and 2018.

As shown in previous research conducted at the Central Bank of Malta, rising female participation rates and the influx of foreign workers had a significant impact on the economy's potential output growth (Grech (2017); Micallef (2018)).

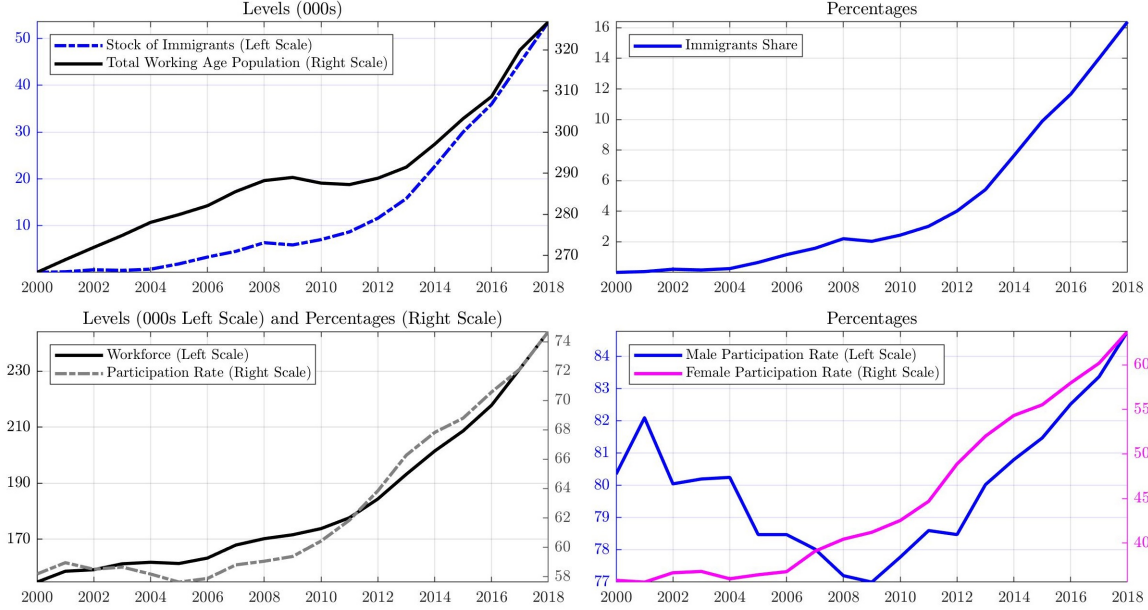


Figure 1: Evolution of: (a) Stock of Immigrants and Working Age Population (Upper Left Graph); (b) Share of Immigrants as a Percentage of the Working Age Population (Upper Right Graph); (c) Workforce and Total Participation Rate (Bottom Left Graph); (d) Male and Female Participation Rate (Bottom Right Graph)

Motivated by the insights taken from figure 1, this paper focuses on the macroeconomic implications of labour supply shocks, both foreign and domestic, in the Maltese economy. The main goal of this work is to provide a structural analysis by running a battery of vector autoregressions that identify both domestic and foreign labour supply shocks and disentangle them from a number of other shocks affecting the Maltese labour market as well as the wider economy.¹ With this aim in mind, this work explicitly and closely follows the empirical strategy implemented in Furlanetto & Robstad (2019).

In terms of impulse responses, the main results of the paper point toward a significant and positive impact of the identified labour supply shocks on domestic production,

¹In this work “foreign labour supply shocks” and “immigration shocks” are used interchangeably.

a downward pressure on wages and unemployment as well as a significant impact on a number of other variables of interest. In terms of forecast error variance decomposition, the two shocks are able to explain a substantial portion of the unexpected fluctuations in gross domestic product, immigrants share in the economy and unemployment rate.

The rest of this work is organised as follows. Section 2 outlines the methodology employed by showing the baseline model and by describing the identification strategy of the shocks of interest. Section 3 presents the results of this study. More precisely, it shows the results of the baseline model as well as those stemming from a number of additional exercises. Finally, section 4 concludes.

2 Methodology

This section outlines the methodology implemented in this paper. More precisely, it shows the reduced form model of the Bayesian VAR to be estimated and provides information regarding the choice of the priors as well as the estimation setup. Finally, and more importantly, this section describes how the shocks of interest are identified.

2.1 Model

The model to be estimated assumes Malta to be an economy whose domestic variables are explained by a number of their own lags plus an exogenous component. More precisely, the reduced form of the baseline model has the following VAR representation:

$$y_t = A + \sum_{l=1}^L B_l y_{t-l} + Ct + u_t \quad (1)$$

for $t = 1, \dots, T$. In 1 y_t is an $N \times 1$ vector of domestic and endogenous variables and y_{t-l} a number of lagged values of the latter with $l = 1, \dots, L$. A is an $N \times 1$ vector of intercepts while B_l represent $N \times N$ matrices containing the slopes relative to the lagged values of the endogenous variables. C is an $N \times 1$ vector of coefficients relative to the linear time trend t and, finally, u_t is an $N \times 1$ vector of reduced form residuals with $u_t \sim N(0, \Sigma)$ where Σ is the variance-covariance matrix.

The N endogenous variables representing the Maltese economy are gross domestic product, real wages, the participation rate series defined as the ratio between labour force (both domestic and foreign) and total working age population (i.e., individuals whose age is between 16 and 64 years), the immigrants share series defined as the stock

of immigrants residing in Malta over the total working age population and, finally, the unemployment rate. Data are collected at quarterly frequency and the sample goes from 2004Q1 to 2019Q2. The sample starts in 2004Q1 because the influx of foreign labour force into the Maltese economy was only a marginal phenomenon before the accession to the European Union. All the variables enter in logarithms with the exception of the unemployment rate that enters in percentage points.² As typically done with such a frequency, the lag length is set to four in order to cover a complete year of data. More details regarding data sources, definitions and transformations can be found in Appendix A.1.

The VAR is estimated with a natural conjugate prior in order to guarantee that the posterior distributions belong to the same family of the priors Kadiyala & Karlsson (1998). To this end, I use a Normal inverse Wishart prior. The model is estimated by setting the total number of Gibbs sampler iterations in such a way to collect 5,000 retained draws from the posterior distributions of the VAR parameters. Such a number is deemed to be satisfactory in order to conduct a meaningful inference. More information regarding the estimation procedure, i.e., choice of the lag length, how the priors are set, likelihood, posterior distributions, convergence diagnostics, can be found in appendices A.2, A.3 and A.4.

2.2 Identification Strategy

In order to meaningfully link the reduced form shocks with the structural ones it is necessary to resort to a reliable identification strategy. In this work, the identification of the model is implemented by closely following the sign restrictions in Furlanetto & Robstad (2019). Such a strategy is a simple and effective way to disentangle three shocks originating in the labour market from other two originating in the wider economy.

The sign restrictions impose restrictions on the decomposition of the estimated variance covariance matrix Σ . In order to achieve this task, the prediction error u_t can be written as a linear combination of the underlying structural innovations v_t , i.e., $u_t = Dv_t$, with the structural shocks being distributed as standard normals, i.e., $v_t \sim N(0, I_N)$ where I_N is a $N \times N$ identity matrix. The objective is to decompose the estimated variance-covariance matrix in a suitable way, i.e., $\Sigma = DD'$, in order to collect a set of models having the entries of D complying with a number of restrictions that are justifiable from a theoretical standpoint. As suggested in Canova & Paustian

²For this reason, the inclusion of a time trend in the VAR aims at improving the estimation of the parameters as all the variables have a trending behaviour.

(2011), the sign restrictions are imposed only on impact and are implemented as shown in table 1:

	<i>Business Cycle</i>	<i>Wage Bargaining</i>	<i>Domestic Labour Supply</i>	<i>Immigration</i>	<i>Residual</i>
<i>Real GDP</i>	+	+	+	+	+
<i>Real Wages</i>	+	-	-	-	+
<i>Participation Rate</i>	+	-	+	+	-
<i>Immigrants Share</i>	?	?	-	+	?
<i>Unemployment</i>	?	?	?	?	?

Note: The signs refers only to the impact responses of the variable of interest $y_{i,t}$ due to a structural shock $v_{j,t}$. A '+' indicates that $\delta y_{i,t}/\delta v_{j,t} > 0$ while a '-' indicates that $\delta y_{i,t}/\delta v_{j,t} < 0$. Finally, a '?' indicates that no restriction is imposed on that variable.

Table 1: Impact Sign Restrictions Characterising the Structural Shocks of Interest

The restrictions aim at fully identifying the model and, as clearly outlined in Furlanetto & Robstad (2019), are mainly derived from a set of theoretical models that study the effects of the five shocks of interest. A generic business cycle shock is defined as a shock that produces an increase in gross domestic product and real wages and, in addition, an increase in the participation rate. Such restrictions are meant to capture foreign shocks and different kinds of demand shocks (e.g., government spending and Euro area monetary policy shocks).

A wage bargaining shock produces an increase in production but also countercyclical dynamics in wages and participation rate. Such restrictions are justified in a New Keynesian model with search and matching frictions and endogenous labour force participation as in Foroni et al. (2018) but also in the model of unemployment outlined in Gali et al. (2011).

The subsequent two shocks are strongly linked to one another and represent the focus of this work. They both represent labour supply shocks and, as such, they both increase production and participation rate while decreasing real wages. As outlined in Furlanetto & Robstad (2019), and corroborated also by a two-country DSGE model with migration in Hart & Clemens (2019), what distinguishes the two is the response of the immigration share. On the one hand, a domestic labour supply shock decreases the immigration share while, on the other hand, an immigration shock generates the opposite effect. Foreign labour supply shocks are meant to capture the recent influx of immigrants while domestic ones are meant to capture the recent surge in the domestic

participation. For both shocks, the assumption of positive co-movement between production and participation rate aims at capturing that portion of labour supply that rapidly enters the labour force, i.e., that portion that finds a job or that is actively looking for one. Furthermore, the assumption of negative impact on real wages finds theoretical support in DSGE models aiming at modelling labour supply shocks, i.e., Gali et al. (2011) and Foroni et al. (2018), and immigration as a form of labour supply shock, i.e., Borjas (2003), Ottaviano & Peri (2008) and Ottaviano & Peri (2012). These models assume that real wages are not affected by immigration shocks in the long run once the adjustment of capital is completed. On the other hand, in the short run, they do affect real wages in a negative way, i.e., as long as the capital adjustment is not complete, and even more on impact as the capital stock is fixed.

Finally, a fifth shock is introduced in order to fully identify the model as, by doing so, the number of endogenous variables will equate the number of identified shocks. The latter is a residual shock and is defined as a shock that increases production and wages but, at the same time, decreases the participation rate on impact. As discussed in Furlanetto & Robstad (2019), these dynamics are generated by technology shocks justifiable in light of Foroni et al. (2018) and Campolmi & Gnocchi (2016).

3 Results

This section shows the main results obtained from the estimation of the model in equation 1, i.e., impulse responses to the two labour supply shocks as well as the forecast error variance decomposition of all the variables involved in the estimation in relation to all the five identified shocks. Furthermore, it shows the effect of the two labour supply shocks on a number of selected variables of interest to the Maltese economy and a sensitivity analysis. The latter aims at assessing the robustness of the baseline results when the assumption of a decrease in real wages on impact is relaxed. Throughout this section, the results relative to the impulse responses will be based on the posterior distribution of the retained responses while the results relative to the forecast error variance decompositions will be obtained by employing the median draw.³

³We are aware of the critique outlined in Fry & Pagan (2011) and the ensuing possibility of using their solution, i.e., the Median Target Method. The latter is correct from a theoretical standpoint but is more difficult to implement from a practical one. It is so as the draw that generates impulse responses that are the closest to the median will likely differ on the basis of many factors, e.g., number of sampler draws, random generator seeds, etc. As such, the results are likely to change from estimation to estimation making difficult any meaningful inference.

3.1 Baseline Results

Figure 2 shows the responses to a foreign labour supply shock in the estimated baseline model. More precisely, it shows the median responses together with the 68% credible bands. The response horizon is set to forty quarters and, given the unit of measure of the variables used, all the responses are expressed in percent while those relative to the unemployment rate are expressed in percentage points. The shock is normalised to generate a median increase in the share of immigrants into the Maltese economy by 1% on impact.

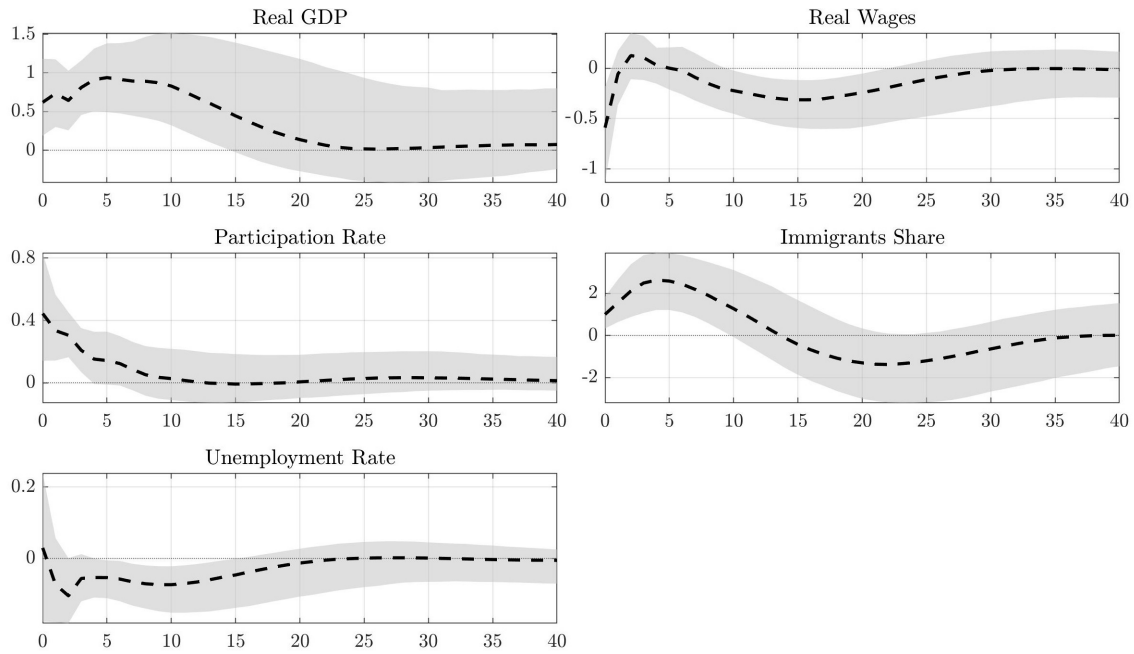


Figure 2: Impulse Responses to a Foreign Labour Supply Shock at each Horizon in the Baseline Model - Posterior Median Values (Black Dashed Lines) and 68% Bands (Grey Shaded Areas)

A foreign labour supply shock associated with a median exogenous 1% increase in the immigrants share makes the participation rate increase by roughly 0.45% on impact and such a response remains significant for four quarters. The response of the immigrants share itself is hump-shaped. More precisely, in median terms, after a 1% increase on impact the response goes up to slightly below 3% before becoming no longer significant after around ten quarters. The hump-shaped response of the

immigrants share can be explained, in a similar way to Furlanetto & Robstad (2019), by late registrations to the immigration authority as well as in light of possible family reunifications even though the latter have historically played a marginal role in the Maltese economy. An immigration shock produces a persistent positive response of the real output accompanied by an overall negative response of real wages. As far as real output is concerned, a 1% foreign labour supply shock increases real gross domestic product by 0.6% on impact and such a response is hump-shaped with a peak response of slightly below 1% after five quarters. The response becomes no longer significant after around fifteen quarters. The response of real wages shows a particular behaviour. More precisely, an immigration shock produces a downward pressure on wages of about 0.5% on impact but then, after going back to zero in the medium term, it produces a smaller downward pressure of around 0.25% that turns insignificant after twenty-two quarters. Finally, an immigration shock significantly lowers the unemployment rate over medium term horizons. Such a decrease can be explained by the fact that generally migrants move to Malta after having already received a job offer. By doing so, by entering the Maltese labour market, they contribute to lower the unemployment rate.

Figure 3 shows the responses to a domestic labour supply shock in the estimated baseline model. In an attempt to make them more comparable with those stemming from an exogenous increase in foreign labour supply, the shock is normalised to decrease the share of immigrants in the Maltese labour force by 1% in median terms on impact.

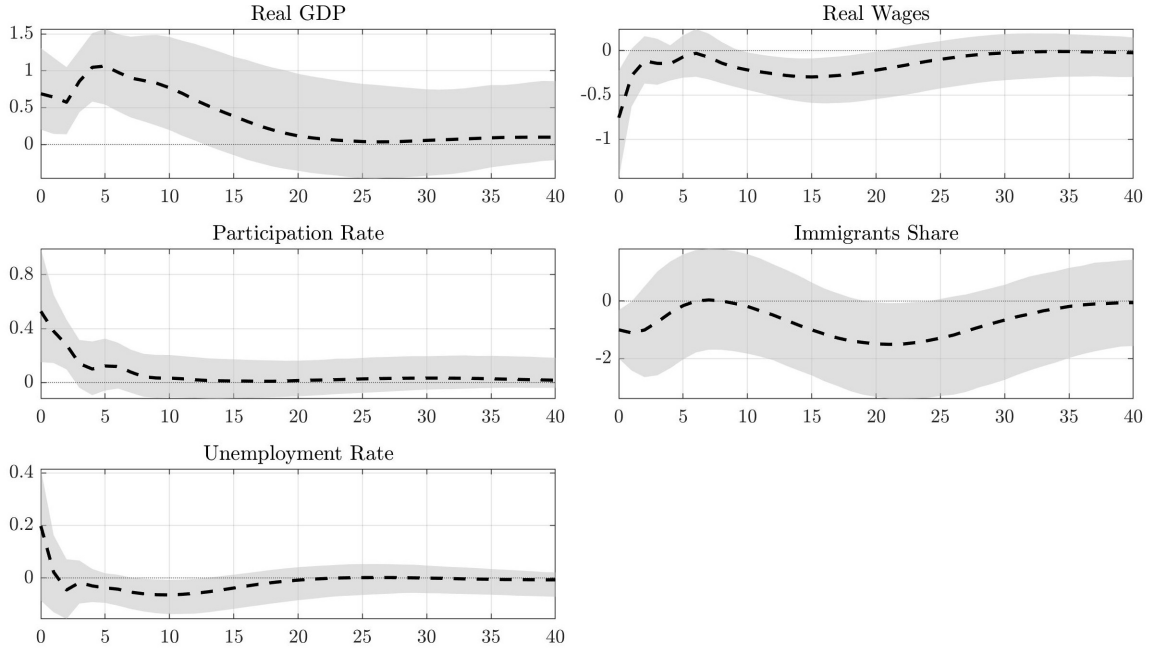


Figure 3: Impulse Responses to a Domestic Labour Supply Shock at each Horizon in the Baseline Model - Posterior Median Values (Black Dashed Lines) and 68% Bands (Grey Shaded Areas)

Figure 3 shows how a domestic labour supply shock produces effects that present small differences with respect to an immigration shock. The shock makes the participation rate increase by roughly 0.5% and such a response is positive and significant for around three quarters. The response of the immigrants share crosses the zero line throughout the entire response horizon thus indicating how a domestic labour supply shock has no significant impact on the immigrants share. If anything, by looking at the median response only, the negative impact is obtained by design of the identification strategy while the negative and persistent response of the immigrants share might reflect a compositional effect on the labour force composition in Malta. The response of real GDP looks very similar to that obtained from an immigration shock. More precisely, the response is hump-shaped and significant for thirteen quarters and the median response is around 0.6% on impact with a peak of slightly more than 1% percent after four quarters. The two most interesting responses are those associated with real wages and unemployment. Even though they look, from a qualitative standpoint, very similar to those stemming from an immigration shock, there are some differences that are worth examining in more detail. A domestic labour supply shock produces a downward

pressure on wages of roughly 0.75% on impact and then presents a negative S-shaped response that hovers -0.25% before turning insignificant after twenty quarters. Such a higher downward pressure might be explained in light of the particular response of unemployment. The latter response is largely positive on impact and equal to $+0.2\%$ in median terms, then it turns negative and, finally, it fades away after around fourteen quarters. The shape of such a response can be explained by the fact that Maltese people who decide to join the labour force are more likely to have the option to wait and see whether they can land a good job, thus putting an initial upward pressure on unemployment. Therefore, this might reconcile with the higher downward pressure on wages detected as firms find themselves with a higher bargaining power when recruiting new workers.

Figure 4, shows the effect of a business cycle and a wage bargaining shock on participation rate, immigrants share and unemployment. The shocks are all normalised to produce a median increase in real GDP by 1% on impact.⁴

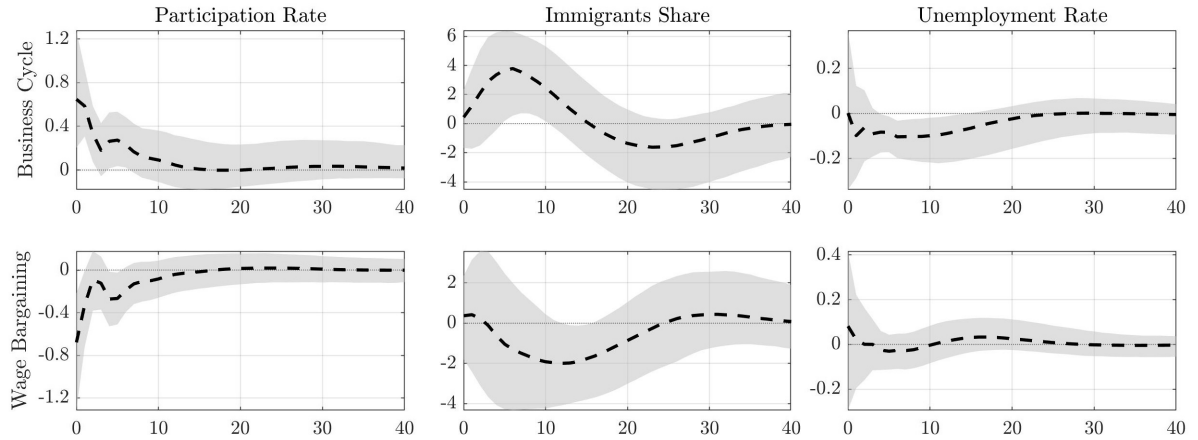


Figure 4: Impulse Responses of Participation Rate, Immigrants Share and Unemployment Rate to a Business Cycle and a Wage Bargaining shock at each Horizon in the Baseline Model - Posterior Median Values (Black Dashed Lines) and 68% Bands (Grey Shaded Areas)

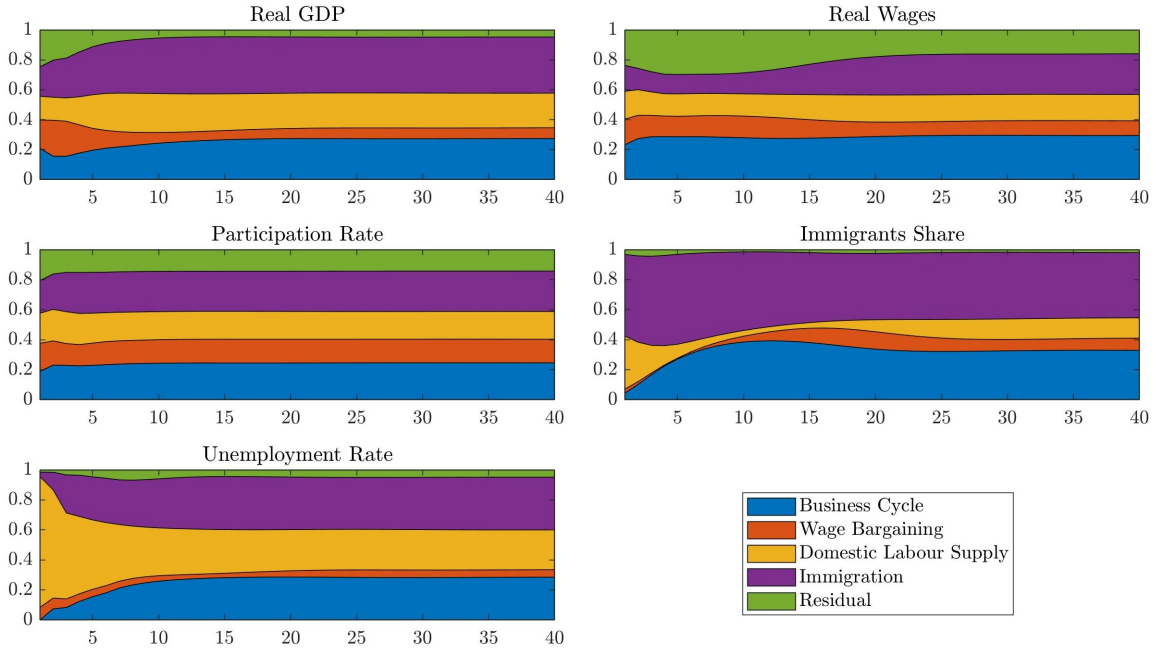
A business cycle shock is accompanied by an increase in the participation rate of

⁴ The responses of all the variables as well as the responses to a residual shock are shown in Appendix A.5.

roughly 0.75% on impact and such a response is significant for around eight quarters. More importantly, the immigration share presents a hump-shaped response. By looking at the median response, the immigrants share increases by 0.5% on impact with at significant peak of around 4% after around six quarters. This sheds lights on how favourable economic conditions induce an influx of foreign workers over the medium term. At longer horizons, this response reverts back to zero. Finally, the unemployment rate presents a significant negative response at peak of nearly 0.1% which takes place after eight quarters.

Finally, a wage bargaining shock discourages the labour force, both domestic and foreign, to enter the market. As such, the impact response of the participation rate is negative and roughly equal to -0.7% on impact while the immigrants share response hovers a significant -2% at peak after twelve quarters. Despite this, the response of unemployment is not significant throughout the entire response horizon.

Figure 5 shows the forecast error variance decomposition of the baseline model. It aims at showing how much of the variability of the unexpected fluctuations of each of the endogenous variables is imputable to each of the identified shocks at several horizons.



Note: The estimation of the forecast error variance decomposition derives from the usage of the median draw.

Figure 5: Forecast Error Variance Decomposition at each Horizon in the Baseline Model

Figure 5 clearly shows how the two identified labour supply shocks contribute in a remarkable way to the unexpected fluctuations of all the endogenous variables. With regard to the real GDP the two shocks explain roughly 35% of its unexpected fluctuations on impact and around 60% at a forty-quarter horizon. More precisely, immigration shocks explain 20% on impact and 40% in the long run, while domestic labour supply shocks respectively around 15% and 20%. The third most important driver of real GDP is the identified business cycle shock which explains around 20% on impact and 35% in the long run.

The importance of labour supply shocks is evident also in explaining the fluctuations of real wages. As the top right panel in figure 5 shows, the two shocks explain around, and roughly in an even way, 40% of the fluctuations of wages at all horizons. The three remaining shocks, i.e., business cycle, wage bargaining and residual, respectively contribute with around 30%, 10% and 20% of the variability.

With regard to the participation rate, the unexpected fluctuations of the latter appear to be explained in fixed proportions by the five identified shocks at all horizons. More precisely, by business cycle shocks for around 25%, wage bargaining shocks for 15%, domestic labour supply shocks for 20%, immigration shocks for 25% and residual shocks for 15%.

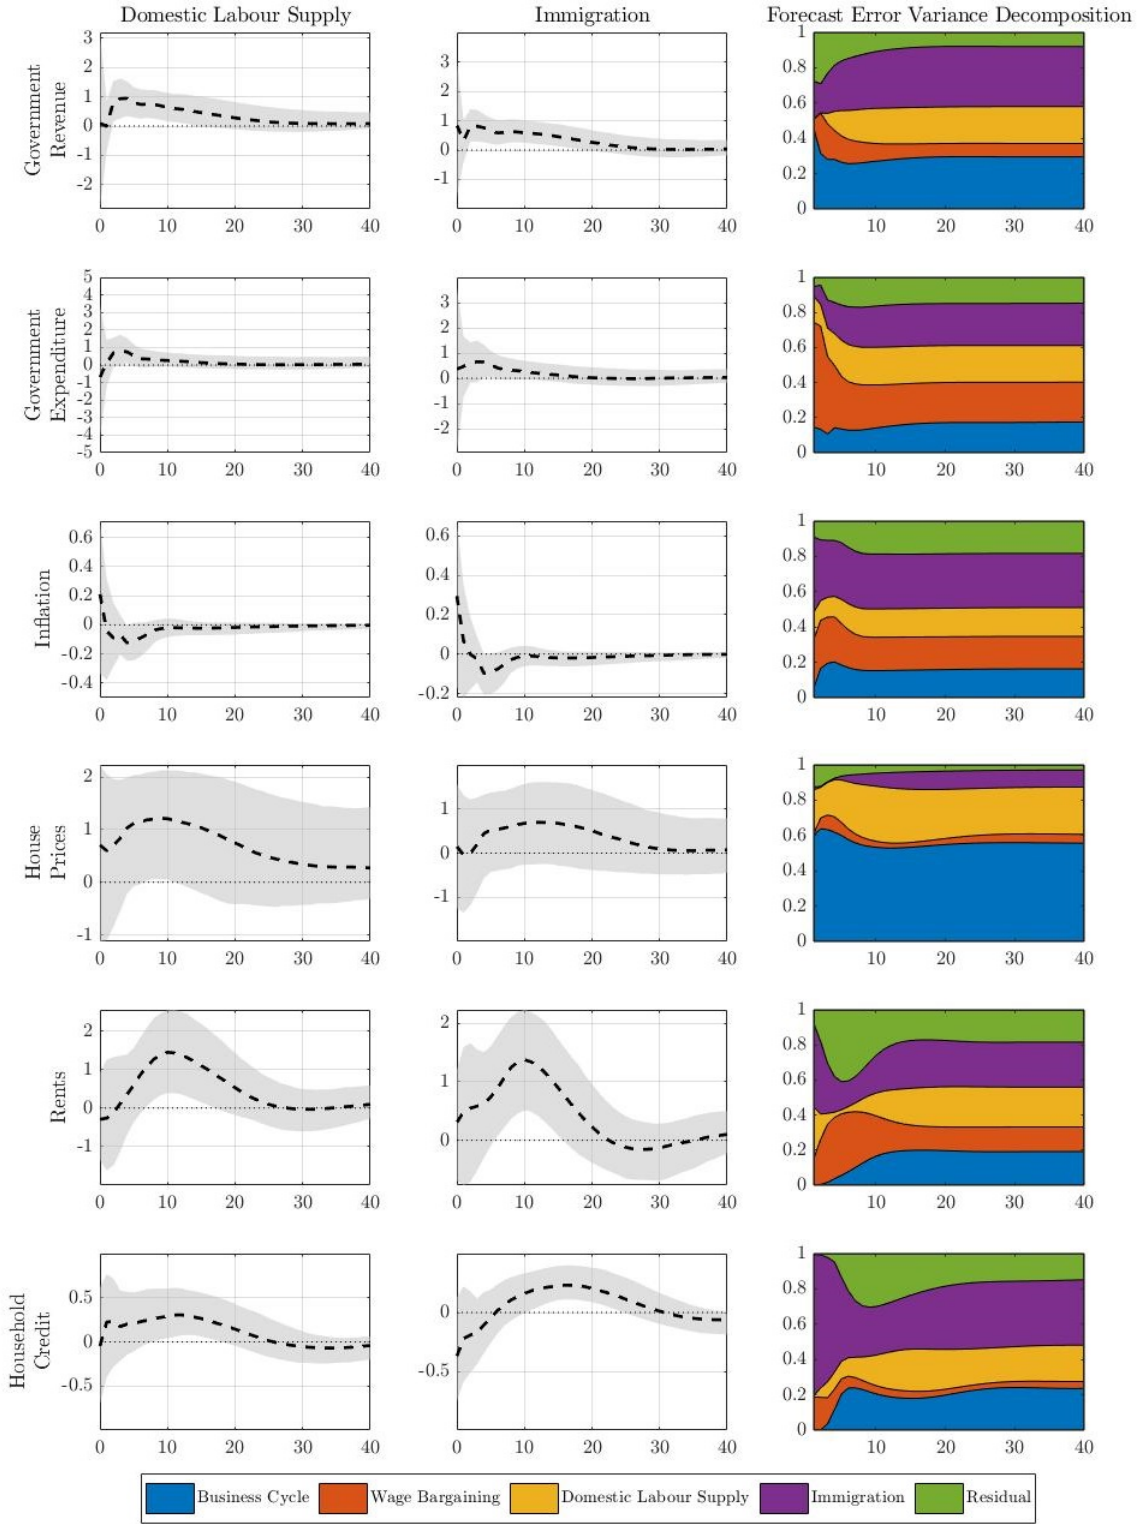
The importance of the two identified labour supply shocks reaches its peak in explaining the two remaining labour related variables, i.e., immigrants share and unemployment rate. The bottom right panel clearly shows how the proportion of immigration in the economy is mainly driven by its own shocks. On impact they account for around 50% of their own fluctuations but this percentage decreases as the horizon grows large and eventually goes down to a value of around 40%. Such a large share of forecast error variance decomposition at long horizons might be explained in light of family reunifications and/or by the fact that migrants tend to go where they find a community belonging to their country of origin. In addition, the contribution of domestic labour supply shocks is remarkable on impact as they explain around 40% of the unexpected fluctuations of the immigrants share. This result can be explained in light of the lack of domestic labour supply with respect to the level that would be justifiable by the economic conditions on the Maltese islands. As the horizon grows large, though, the contribution of domestic labour supply shocks gets smaller and stabilises at around 15%. The third main driver of the fluctuations of the immigrants share is represented by business cycle shocks. The latter, even though contributing with a little 5% on impact, explain more than 30% of the immigrants share fluctuations in the long run.

The remaining two shocks do not appear to have a relevant role.

Finally, and straightforwardly, the main drivers of unemployment are the two labour supply shocks as well as the business cycle one. In particular, the fluctuations of unemployment are explained by domestic labour supply shocks for 85% on impact and 25% at long horizons. It follows that immigration and business cycle shocks play a smaller role on impact, around 5% for the former and negligible for the latter, but their relevance increases as the horizon grows large. At a forty-quarter horizon, foreign labour supply shocks explain around 35% of unemployment fluctuations while business cycle shocks contribute with around 30%. Similarly to the case of the immigrants share, wage bargaining and residual shocks do not appear to play a relevant role.

3.2 Further Experiments on Selected Variables of Interest

This subsection presents some experiments in order to assess the relevance of the identified labour supply shocks on nine selected variables of interest. Each experiment is conducted by utilising the baseline model in 1 and then by substituting the unemployment rate series with one variable of interest at a time. In such a way, the model is still fully identified as the number of endogenous variables is equal to the number of identified shocks. Figure 6, for the sake of space, shows the results related to the first six of the selected variables. On the first and second columns the impulse responses of only the variables of interest are shown, while their associated forecast error variance decompositions are reported on the rightmost column. As in the previous subsection, the two shocks are normalised to increase or decrease the immigrants share by 1% on impact. As all the variables of interest of these additional exercises enter in logarithms, the responses are expressed in percent.



Note: The estimation of the forecast error variance decompositions derive from the usage of the median draw.

Figure 6: Impulse Responses to a Domestic and a Foreign Labour Supply Shock at each Horizon on Selected Variables of Interest - Posterior Median Values (Black Dashed Lines) and 68% Bands (Grey Shaded Areas) (First and Second Columns) and Forecast Error Variance Decomposition (Third Column)

Public Finances

The first two rows of figure 6 aim at shading some light on the burden that immigrants place on public finances. The latter has always been a cause of ignited debate among policy makers and has always acted as a divider between those in favour and those against the influx of foreign workers. The results suggest how the two labour supply shocks play a similar role in influencing the government variables, i.e., revenues and expenditure. More precisely, both shocks produce responses of government revenues that have a not significant effect on impact but that reach a significant peak of around +1% after three quarters. This can be explained in light of the higher tax revenue the government collects as soon as new workers enter the labour force. At long horizons, immigration shocks explain roughly 30% of the unexpected fluctuations in government revenues while domestic labour supply shocks around 20%. In addition, business cycle shocks explain around 30%.

The responses of government expenditure to the two shocks appear to have a similar shape of those relative to the previous variable but are not as significant in the medium term. The impact responses are not significant but reach a median peak of slightly below 1% after around three quarters. The increase in government expenditure might reflect the higher financial burden implied by family friendly measure aiming at increasing the labour participation. More precisely, this can be explained in light of the necessity to expand social security programs to the new arrived (Borjas (1999)), e.g., the free childcare system, and it might also reflect the necessity for the government to invest in public goods projects. As for the forecast error variance decomposition, the two labour supply shocks explain around, and in an even way, 40% of the fluctuations of the government expenditure. Finally, besides the business cycle shock explaining around 20% of the fluctuations, it is important to notice the relevance of the wage bargaining shock. The latter explains around 60% of the fluctuations on impact and slightly more than 20% at longer horizons. This probably reflects how a large part of the government expenditure is represented by wages in the public sector and therefore how wage negotiations by labour unions play a relevant role. All in all, it is possible to argue that the impact of labour supply shocks on public finances is positive as the responses of revenue appear to offset those of expenditure.

Inflation

The response of inflation is insignificant throughout the entire response horizon for both shocks but, by looking at the median responses, it is positive on impact. More precisely, a domestic labour supply shock increases inflation by 0.2% while an immigration shock by 0.3%. The higher effect of the latter might be due to the expenses that immigrants are likely to incur when moving for work reasons. As for its drivers, business cycle unexpected movements explain around 20% of the unexpected fluctuations in inflation, domestic labour supply around 15% while immigration around 35%.

Real Estate

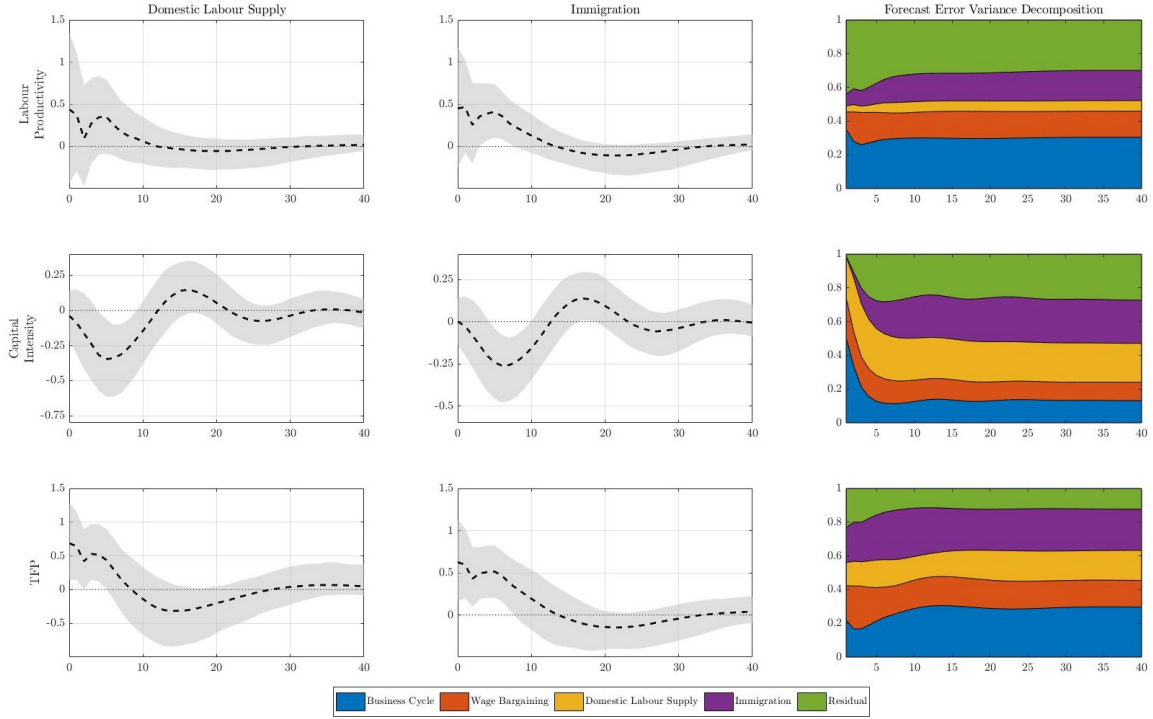
Besides the financial burden posed on government finances, the response of real estate variables is another hot topic among policy makers when discussing the role of immigration. In line with recent literature (Saiz (2003), McDonald (2013) and Dalbis et al. (2019)), the fourth and the fifth rows of figure 6 try to answer some questions. Both labour supply shocks produce hump-shaped responses of house prices that are not significant throughout the entire response horizon. However, if we look at the median responses, house prices remain slightly more than 1% higher for around 20 quarters after a domestic labour supply shock hits the economy, while after an immigration one slightly less than 1%. For the latter, the impact response is almost null. As for the forecast error variance decomposition, the most relevant role is played by unexpected movements in the business cycle for around 50% followed by domestic labour supply with around 30% and immigration with around 10%. The bigger contribution of domestic labour supply shocks relatively to foreign ones might be due to the higher desire of Maltese people to buy a property as soon as they enter the labour force.

The second real estate-related variable of interest is rents. Both shocks produce similar hump-shaped responses that significantly reach a median peak of roughly 1.5% after ten quarters. Even though the impact responses cross the zero line, in median terms the domestic shock generates a slightly negative response while, as expected, a foreign one immediately contributes in a positive way by nearly 0.3%. The forecast error variance decomposition shows how the contribution of the identified shocks is likely to vary with the horizon. In the long run, the fluctuations of rents are mainly explained by immigration, domestic labour supply and business cycle shocks respectively for 25%, 20% and 20%. The most interesting results are at short horizons though. On impact immigration shocks are the main contributors with around 45%. This is not surprisingly as immigrants are in all likelihood those who need an accommodation

to rent. The second most important contributor on impact is domestic labour supply shocks with around 25% thus indicating how even Maltese people might need to rent a property as soon as they enter the labour market. This might reflect the desire to live independently and or a possible desire to relocate within the Maltese territory. Finally, wage bargaining shocks play a relevant contribution with slightly less than 20% on impact and around 40% after five quarters. This result sheds light on how wage negotiations have a relevant impact on rents as they help determine the willingness to pay for accommodation.

Household Credit

Finally, the responses of household credit per capita vary substantially between the two shocks. On the one hand, for the domestic case household credit presents an insignificant response on impact that then turns significantly positive peaking at 0.25% in median terms after twelve quarters. This might reflect the increase in the stock of mortgage loans to buy a property as soon Maltese people enter the labour force. On the other hand, by looking at the median responses to an immigration shock, the per capita household credit goes down by 0.4% in all likelihood as a result of the population increase. As for its drivers, the unexpected fluctuations of household credit per capita are explained by immigration shocks for around 80% on impact and 40% at long horizons. Domestic labour supply and business cycle shocks explain very little on impact but then they both contribute for 20% each at long horizons.



Note: The estimation of the forecast error variance decompositions derive from the usage of the median draw.

Figure 7: Impulse Responses to a Domestic and a Foreign Labour Supply Shock at each Horizon on Selected Variables of Interest - Posterior Median Values (Black Dashed Lines) and 68% Bands (Grey Shaded Areas) (First and Second Columns) and Forecast Error Variance Decomposition (Third Column)

Productivity

The last three variables of this sections are relative to the effect of labour supply shocks on three measures of productivity. The first row of figure 7 focuses on labour productivity. The responses to the two shocks are very similar showing a median 0.5% increase on impact. The main difference is that the response to a domestic labour supply shock is not significant at any horizon while that to an immigration one is mostly significant for eight quarters. The two shocks explain roughly 20% of the unexpected fluctuation of labour productivity while the main drivers of the latter are mostly business cycle and residual shocks with around 30% each.

The response of capital intensity is another important topic as this may help shed light on the degree of substitutability and/or complementarity between capital and labour as soon as a labour supply shock hits the economy. The responses to the two

shocks are roughly the same. They present a hump-shaped behaviour with a negative and significant peak of -0.25% after around six quarters. Afterwards, the response reverts back and hovers the zero line. The negative and similar responses of capital intensity suggests how after these shocks hit the economy, Maltese firms are more likely to employ more labour for any given level of capital. The capital intensity ratio is, at long horizon explained by the five identified shocks for around, respectively in the order, 15%, 10%, 25%, 25% and 25%. What is important to notice is, however, how the impact of business cycle shocks is relevant on impact as its contribution jump to 50% of the unexpected fluctuations. A possible argument for this is that changes in economic conditions immediately influence firms' decisions regarding their capital-labour composition.

Finally, the responses of total factor productivity are similar across labour supply shocks. In both cases, the median impact response is significantly positive at around 0.6%. For the domestic case the response becomes insignificant after roughly one year, while it is more persistent in the foreign one. The main difference between the two is the more pronounced negative adjustment taking place in the medium term after a domestic labour supply shock. As suggested in Peri & Sparber (2009) and Peri (2016), in countries with large inflows of immigrants, natives with a lower education tend to specialise in more communication-intensive tasks thus leaving more manual-intensive tasks to immigrants. Such a rebalancing process would result in efficiency gains for the economy as a whole. Finally, the unexpected fluctuations of this measure of productivity have the business cycle as the main contributor with around 30% while the two labour supply shocks contribute for 40% (15% for the domestic and 25% for the foreign). The remaining two shocks explain the remain 30% of the fluctuation of total factor productivity.

To summarise, both labour supply shocks exerted a positive impact on labour productivity through an increase in total factor productivity, whereas the impact on the capital intensity is negative. The increase in productivity, in line with Micallef & Ellul (2017), explains the importance of the increase in labour supply in addressing crucial skills and labour shortages, which were key to raise the country's supply-side potential.

3.3 Sensitivity

The aim of this subsection is to assess the robustness of the baseline results to a change in the identification strategy. The focus is on the response of real wages by

relaxing the assumption that the latter have to fall on impact as soon as the labour supply shocks hit the economy. More precisely, the assumption of negative response of wages is replaced with a zero restriction implying that they do not move on impact. As it is widely argued in the literature (Woodford (2003), Levin et al. (2005), Taylor (2007)), wages might not adjust quickly to changes in labour market conditions such as when labour supply shocks hit an economy. This slow moving behaviour is mainly due to the fact that salaries are sticky because are often predetermined as they are the result of labour unions negotiations that take place only in certain points in time. In addition, and as previously argued, for the case of migration the impact crucially depends on the skill composition of migrants and natives. Figures 8 and 9 present the responses to the two shocks that, as previously done, imply a median 1% increase or decrease in the immigrants share on impact. The grey shaded areas represent the responses under the baseline identification while the red lines those under this modified one.

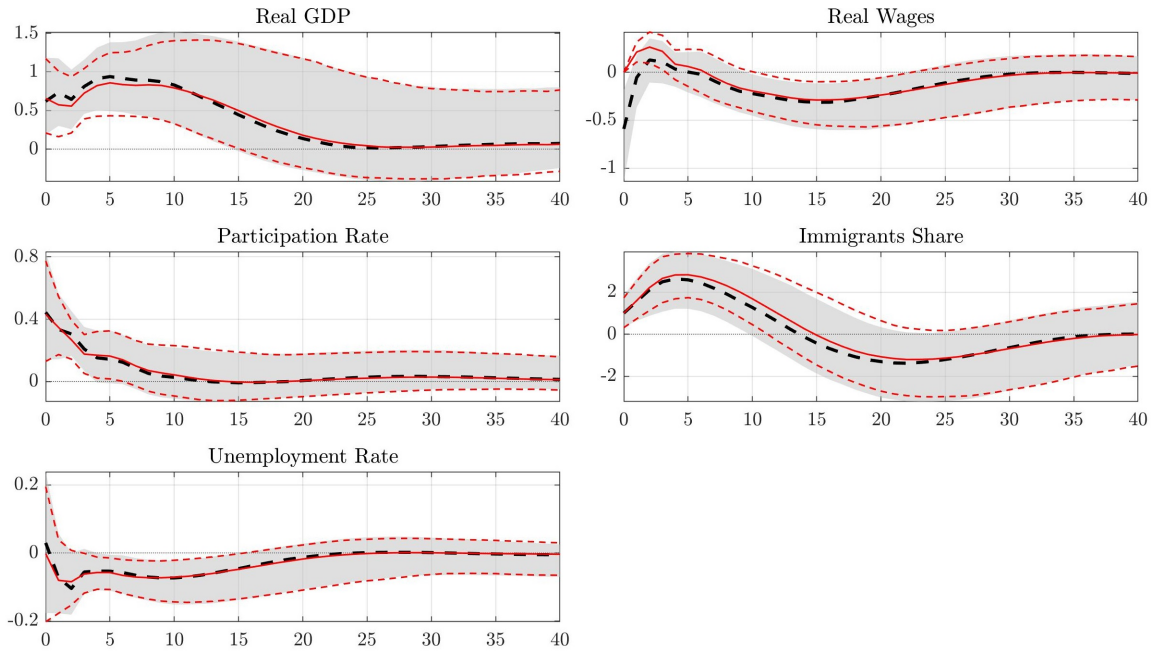


Figure 8: Impulse Responses to a Foreign Labour Supply Shock at each Horizon with the Baseline Identification Strategy (Black Dashed Lines and Grey Shaded Areas) and Under the Assumption of Sticky Wages (Red Lines) - Posterior Median Values and 68% Bands

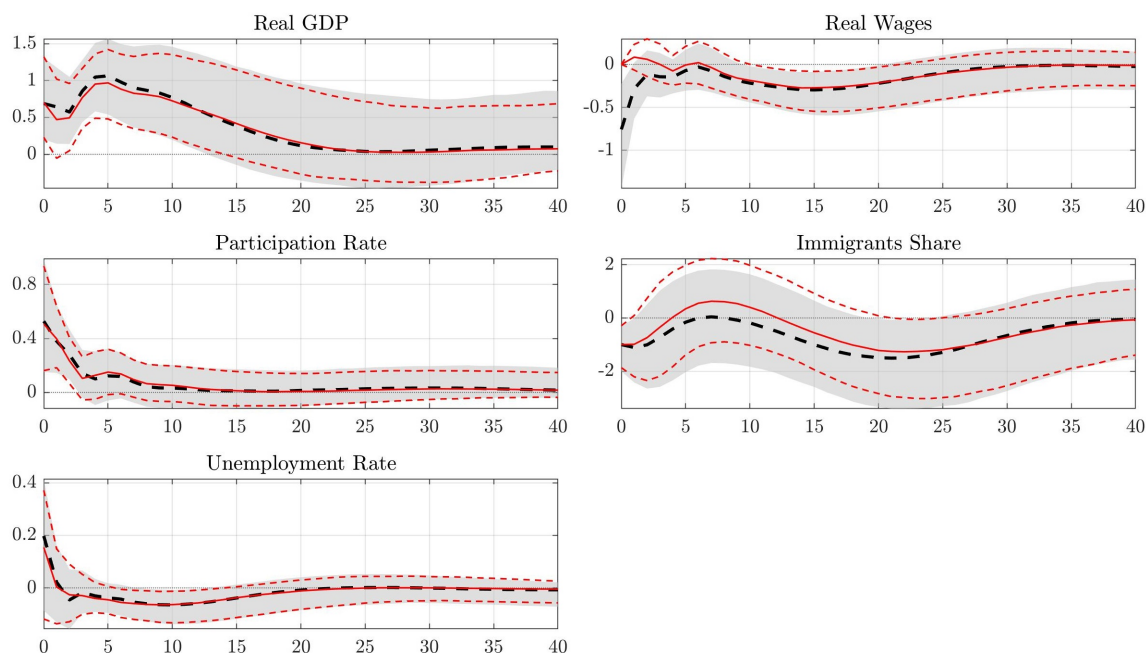


Figure 9: Impulse Responses to a Domestic Labour Supply Shock at each Horizon with the Baseline Identification Strategy (Black Dashed Lines and Grey Shaded Areas) and Under the Assumption of Sticky Wages (Red Lines) - Posterior Median Values and 68% Bands

Figures 8 and 9 demonstrate how the assumption of sticky wages is not relevant in tracing the medium and long term dynamic responses of real wages after the two labour supply shocks hit the Maltese economy. It is so as the credible regions generated by the two identification strategies overlap almost perfectly with the obvious exception of the impact and the short term responses. The results suggest that real wages might also increase in the short-run following an immigration shock, in line with the evidence from the micro literature, though the impact over the medium-term is quantitatively and qualitatively similar to the baseline model. A similar finding is obtained with the domestic labour supply shock, though the short-run effect on wages is not different from zero.

4 Conclusions

Over the recent years Malta has experienced a remarkable increase in its labour force due to a large influx of immigrants and an unprecedented increase in its domestic participation. The observation of this relatively new phenomenon has posed some questions

regarding the impact that such an increase in the labour force, as well as its composition, is likely to exert on the Maltese economy. In order to address these questions, this work focused on the macroeconomic implications of labour supply shocks on the Maltese economy by disentangling them from other shocks originating in the wider economy. The analysis is conducted by estimating a baseline structural vector autoregression as well as a number of additional ones aiming at understanding how both domestic and foreign labour supply shocks affect a number of selected macroeconomic variables of interest. All the structural VARs are estimated by means of Bayesian techniques while the shock identification strategy is implemented through sign restrictions as described in Furlanetto & Robstad (2019). The sample of data covers the period in which Malta has been part of the European Union, i.e., from 2004Q1 to 2019Q2, as prior to the EU accession foreign workers inflows were only a marginal phenomenon.

The results point towards a relevant impact of both labour supply shocks. For a better comparability between the two, they were normalised to increase (in the case foreign labour supply shocks) or decrease (in the case of domestic labour supply shocks) the share of immigrants in the economy by 1% on impact. Both shocks produce a significant increase in real output with a peak of 1% after around one year, a downward pressure on wage and a significant decrease in unemployment with a peak of about -0.1% after two years and a half. The forecast error variance decomposition confirms how the two labour supply shocks are relevant contributors of the unexpected fluctuations of all the variables in the baseline model. This is particularly relevant for real gross domestic product, immigrants share and unemployment rate as the two shocks combined are able to explain around 60% of their fluctuations in the long run. In addition, broadly in line with the evidence from the recent literature MAC (2018), the identified labour supply shocks contribute to the fluctuations of several other macroeconomic variables of interest, such as the government revenue and expenditure, rents and a number of relevant measures of productivity. Finally, a sensitivity analysis aiming at explicitly considering the presence of sticky wages suggests that real wages might also increase in the short-run following an immigration shock. A similar finding is also obtained in response to a domestic labour supply shock, though the short-run effect is not different from zero.

References

- Arias, J. E., Rubio-Ramirez, J. F. & Waggoner, D. F. (2018), ‘Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications’, *Econometrica* **86**(2), 685–720.
- Blake, A. & Mumtaz, H. (2017), ‘Applied bayesian econometrics for central bankers’, *Centre for Central Banking Studies* .
- Borjas, G. J. (1999), ‘Immigration and welfare magnets’, *Journal of Labor Economics* **17**, 607–637.
- Borjas, G. J. (2003), ‘The Labor Demand Curve is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market’, *The Quarterly Journal of Economics* **118**(4), 1335–1374.
- Campolmi, A. & Gnocchi, S. (2016), ‘Labor market participation, unemployment and monetary policy’, *Journal of Monetary Economics* **79**(C), 17–29.
- Canova, F. & Paustian, M. (2011), ‘Business cycle measurement with some theory’, *Journal of Monetary Economics* **58**(4), 345–361.
- DAlbis, H., Boubtane, E. & Coulibaly, D. (2019), ‘International migration and regional housing markets: Evidence from france’, *International Regional Science Review* **42**(2), 147–180.
- Faroni, C., Furlanetto, F. & Lepetit, A. (2018), ‘Labor supply factors and economic fluctuations’, *International Economic Review* **59**(3), 1491–1510.
- Fry, R. & Pagan, A. (2011), ‘Sign restrictions in structural vector autoregressions: A critical review’, *Journal of Economic Literature* **49**(4), 938–60.
- Furlanetto, F. & Robstad, O. (2019), ‘Immigration and the macroeconomy: some new empirical evidence’, *Review of Economic Dynamics* **34**, 1–19.
- Gali, J., Smets, F. & Wouters, R. (2011), ‘Unemployment in an estimated new keynesian model’, *NBER Macroeconomics Annual 2011* **26**, 329–360.
- Grech, A. G. (2017), ‘Did malta s accession to the eu raise its potential growth? a focus on the foreign workforce’, *Journal of Economic Integration* **32**(4), 873–890.

- Hart, J. & Clemens, M. (2019), ‘A search and matching approach to business-cycle migration in the euro area’, *Annual Conference 2019 (Leipzig): 30 Years after the Fall of the Berlin Wall - Democracy and Market Economy* .
- Kadiyala, R. K. & Karlsson, S. (1998), ‘Numerical methods for estimation and inference in bayesian var-models’, *Journal of Applied Econometrics* .
- Levin, A. T., Onatski, A., Williams, J. C. & Williams, N. (2005), ‘Monetary policy under uncertainty in micro-founded macroeconomic models’, (11523).
- Litterman, R. B. (1986), ‘Forecasting with bayesian vector autoregressions: Five years of experience’, *Journal of Business & Economic Statistics* **4**(1), 25–38.
- MAC (2018), ‘Eea migration in the uk: Final report’, *Migration Advisory Committee publication* .
- McDonald, C. (2013), ‘Migration and the housing market’, *Reserve Bank of New Zealand Analytical Notes AN2013/10* .
- Micallef, B. (2018), ‘Estimating the impact of structural reforms to increase the female participation rate in malta’, *International Journal of Social Science Studies* **6**(8), 73–84.
- Micallef, B. & Debono, N. (2020), ‘The rental sector and the housing block in stream’, *Central Bank of Malta Policy Working Papers Series* .
- Micallef, B. & Ellul, R. (2017), ‘Medium term estimates of potential output growth in malta’, *Challenges and opportunities of sustainable economic growth: the case of Malta - Central Bank of Malta publications* pp. 15–28.
- Ottaviano, G. I. P. & Peri, G. (2008), ‘Immigration and national wages: Clarifying the theory and the empirics’, *NBER Working Papers Series* (14188).
- Ottaviano, G. I. P. & Peri, G. (2012), ‘Rethinking the effect of immigration on wages’, *Journal of the European Economic Association* **10**(1), 152–197.
- Peri, G. (2016), ‘Immigrants, productivity and labour markets’, *Journal of Economic Perspectives* **30**(4), 3–30.
- Peri, G. & Sparber, C. (2009), ‘Task specialization, immigration, and wages’, *American Economic Journal: Applied Economics* **1**(3), 135–69.

- Primiceri, G. E. (2005), ‘Time varying structural vector autoregressions and monetary policy’, *The Review of Economic Studies* **72**(3), 821–852.
- Saiz, A. (2003), ‘Room in the kitchen for the melting pot: Immigration and rental prices’, *Review of Economics and Statistics* **85**(3), 502–521.
- Spiegelhalter, D., Best, N., Carlin, B. & Linde, A. (2002), ‘Bayesian measures of model complexity and fit (with discussion)’, *Journal of the Royal Statistical Society, Series B* **64**, 1–34.
- Taylor, J. B. (2007), ‘Thirty-five years of model building for monetary policy evaluation: Breakthroughs, dark ages, and a renaissance’, *Journal of Money, Credit and Banking* **39**(s1), 193–201.
- Woodford, M. (2003), ‘Interest and prices’, *Princeton University Press* .

A Appendix

A.1 Data Sources and Transformations

This appendix describes the data series and transformations used in this work. All the data series are collected at quarterly frequency with the exception of the stock of immigrants in Malta which is available only on annual basis. The latter, before being used for any transformation, has been transformed into quarterly frequency through a cubic spline interpolation. All the variables in the estimation are seasonally adjusted and enter in logarithms with the exception of the unemployment rate that enters in levels.

Real gross domestic product: real gross domestic product (National Statistics Office).

Real wages: real hourly compensation for employees calculated as total compensation for employees (National Statistics Office) deflated by the inflation rate that excludes energy and food (see “Inflation” below) and divided by the total number of hours worked (National Statistics Office).

Participation rate: participation rate calculated as total labour force, i.e., working people plus people looking for a job (both local and foreign) divided by working age population, i.e., population between 16 and 64 of age (Eurostat - Labour Force Survey).

Immigrants share: stock of immigrants in the Maltese economy (JobsPlus) divided by working age population (Eurostat - Labour Force Survey).

Unemployment rate: total unemployment rate (National Statistics Office).

Government revenue: nominal government revenues (General Government data).

Government expenditure: nominal government expenditure (General Government data).

Inflation: harmonised index of consumer prices (HICP) that excludes energy and food (Eurostat).

House prices: nominal house prices (Central Bank of Malta) deflated by the inflation rate that excludes energy and food.

Rents: Rents series developed in Micallef & Debono (2020)

Household credit: nominal household credit (Central Bank of Malta) deflated by the inflation rate that excludes energy and food and divided by population.

Labour Productivity: real gross domestic product divided by the total number of hours worked.

Capital intensity: ratio between capital stock and labour force (Central Bank of Malta).

Total factor productivity: total factor productivity (Central Bank of Malta).

A.2 Estimation Details

This appendix aims at illustrating the technical details relative to the procedure to estimate the model in section 2. The VAR described in 1 can be compactly written as:

$$\mathbf{Y} = \mathbf{XB} + \mathbf{U} \tag{2}$$

where $\mathbf{Y} = [y_1, \dots, y_T]'$, $\mathbf{B} = [A, B_1, \dots, B_L, C]'$, $\mathbf{U} = [u_1, \dots, u_T]'$ and

$$\mathbf{X} = \begin{bmatrix} 1 & y'_0 & \dots & y'_{-L} & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & y'_{T-1} & \dots & y'_{T-L} & T \end{bmatrix}.$$

For ease of estimation, the model in 2 can be written in vectorised form as

$$\mathbf{y} = (I_N \otimes \mathbf{X})\beta + \mathbf{u} \quad (3)$$

where $\mathbf{y} = \text{vec}(\mathbf{Y})$, $\beta = \text{vec}(\mathbf{B})$ and $\mathbf{u} = \text{vec}(\mathbf{U})$. In equation 3, $\text{vec}()$ is the columnwise vectorisation operator and $\mathbf{u} \sim N(0, \Sigma \otimes I_T)$. The likelihood function in \mathbf{B} and Σ implied by equation 3 is derived in Kadiyala & Karlsson (1998) and assumes the following form:

$$L(\mathbf{B}, \Sigma) \propto |\Sigma|^{-\frac{T}{2}} \exp\left(-\frac{1}{2}(\beta - \hat{\beta})'(\Sigma^{-1} \otimes \mathbf{X}'\mathbf{X})(\beta - \hat{\beta})\right) \exp\left(-\frac{1}{2}\text{tr}(\Sigma^{-1}\hat{S})\right) \quad (4)$$

where $\hat{S} = (\mathbf{Y} - \mathbf{X}\hat{\mathbf{B}})'(\mathbf{Y} - \mathbf{X}\hat{\mathbf{B}})$ and $\hat{\beta} = \text{vec}(\hat{\mathbf{B}})$ with $\hat{\mathbf{B}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$.

We use a normal-inverse Wishart prior for β and Σ as follows:

$$p(\beta, \Sigma) \sim N(\bar{\beta}, \Sigma \otimes \bar{H}) \quad (5)$$

$$p(\Sigma) \sim IW(\bar{S}, \bar{\alpha}) \quad (6)$$

In 5, the prior mean $\bar{\beta}$ is set in such a way to incorporate the belief that each variable in y_t follows a random walk process like in a Minnesota prior as in Litterman (1986). As the variables in the estimation enter in logarithms, this reflects the belief that the latter are characterised by high persistence. As such, considering that $l = 1, \dots, L$ denotes the lag order while $i, j = 1, \dots, n$ the equation number in the vector autoregression, the elements of $\bar{\beta}$ in 5 are set in the following way:

- $\overline{b_{l,i,j}} = 1$ if $l = 1$ and $i = j$, i.e., for the coefficient on first lagged value of the n^{th} endogenous variable in the n^{th} equation;
- $\overline{b_{l,i,j}} = 0$ otherwise.

Furthermore, following Blake & Mumtaz (2017), \bar{H} is a $K \times K$ diagonal matrix with

$K = (NL + 2)$ and diagonal elements set as follows:

- $(\frac{\lambda_0 \lambda_1}{l^{\lambda_3} \sigma_n})$ for the coefficients on the lagged values of the n^{th} endogenous variable in the n^{th} equation;
- $(\lambda_0 \lambda_4)^2$ for all the other coefficients, i.e., intercept and time trend.

The hyperparameters in the \bar{H} matrix in 5 are set in a quite standard way and have the following function:

- λ_0 controls the overall tightness of the prior and is set to 1;
- λ_1 controls the tightness on the coefficients associated with the first lagged values of the endogenous variables and is set to 0.5;
- λ_3 controls the degree of shrinkage on higher order lags of the endogenous variables and is set to 1;
- λ_4 controls the degree of shrinkage on all the other coefficients and is set to 10^4 , i.e., a flat prior on all the exogenous component.

In 6, $\bar{\alpha} = N + 1$ while \bar{S} is an $N \times N$ diagonal matrix with elements equal to $(\frac{\sigma_n}{\lambda_0})^2$. The choice of $\lambda_0 = 1$ implies that the prior scale \bar{S} of the variance-covariance matrix Σ is obtained through the AR(1) residual variances of the N endogenous variables.

We can then obtain the posterior, as the product of a normal distribution for β conditional on Σ and an inverse Wishart distribution for Σ , through the following expression:

$$p(\mathbf{B}, \Sigma | \mathbf{Y}, \mathbf{X}) \propto |\Sigma|^{-\frac{K}{2}} \exp(-\frac{1}{2}(\beta - \tilde{\beta})'(\Sigma^{-1} \otimes \tilde{H})(\beta - \tilde{\beta})) |\Sigma|^{-\frac{T+\tilde{\alpha}}{2}} \exp(-\frac{1}{2}tr(\Sigma^{-1} \tilde{S})) \quad (7)$$

where $\tilde{\alpha} = T + \bar{\alpha}$, $\tilde{H} = (\mathbf{X}'\mathbf{X} + \bar{H}^{-1})^{-1}$, $\tilde{\mathbf{B}} = \tilde{H}(\mathbf{X}'\mathbf{Y} + \bar{H}^{-1}\bar{\mathbf{B}})$ with $vec(\tilde{\mathbf{B}}) = \tilde{\beta}$ and, finally, $\tilde{S} = \mathbf{Y}'\mathbf{Y} + \bar{S} + \tilde{\mathbf{B}}'\bar{H}^{-1}\bar{\mathbf{B}} - \tilde{\mathbf{B}}'\tilde{H}^{-1}\tilde{\mathbf{B}}$. As such, our inference is carried out by drawing from the following posterior distributions:

$$p(\beta, |\Sigma, \mathbf{Y}, \mathbf{X}) \sim N(\tilde{\beta}, \Sigma \otimes \tilde{H}) \quad (8)$$

$$p(\Sigma|\mathbf{Y}, \mathbf{X}) \sim IW(\tilde{S}, \tilde{\alpha}) \quad (9)$$

Finally, as specified in section 2, in order to obtain the shocks identification through a mix of zero and sign restrictions, the procedure described in Arias et al. (2018) is implemented.

A.3 Lag Selection

This appendix aims at justifying the choice of the lag length used to estimate the model in equation 1. I use the deviance information criterion suggested in Spiegelhalter et al. (2002) to gauge the goodness of fit. The measure is calculated by estimating the model in 1 with lag lengths that go from one to five, i.e, from the minimum possible to a lag length covering one year of data plus one quarter. I also consider two specifications of the model: with and without a linear time trend. The values are shown in table 2:

<i>Lag Length</i>	<i>Without Time Trend</i>	<i>With Time Trend</i>
1	-1281.4	-1287.4
2	-1388.5	-1386.1
3	-1405.9	-1414.1
4	-1398.9	-1426.2
5	-1405.6	-1426.8

Table 2: Deviance Information Criteria for Several Specifications of the Baseline Model

The second column of table 2 shows how the model with a time trend should be preferred as it is associated with lower values of DICs for any number of lags. In addition, the third column shows how a model with more lags should be preferred as it is associated with even lower values of the used measure of goodness of fit. I decide to estimate the model by employing four lags as the DICs do not appear to remarkably

change when moving from a model with four to a model with five lags. Using four lags helps contain the number of parameters to be estimated and, at the same time, the information included in one year of data in is still retained.

A.4 Convergence Diagnostics

This appendix shows the convergence diagnostics relative to the baseline model presented in section 2. In order to assess the satisfactory performance of the algorithm and, therefore, its suitability for conducting inference, the *20th* order autocorrelation of the retained draws is employed. To this end, figure 10 shows such values by distinguishing between the slope coefficients associated with the lagged values of the endogenous variables and the draws of the variance-covariance matrix.

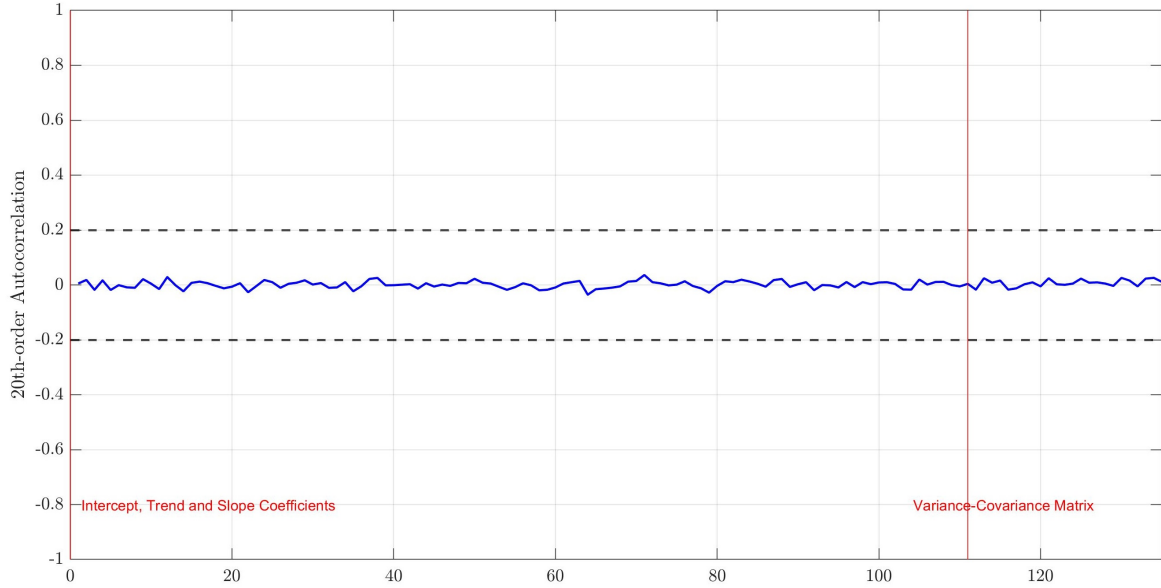


Figure 10: *20th* Order Autocorrelation of the Retained Draws in the Baseline Model

The picture clearly shows how the autocorrelations of the retained draws lie within the $[-0.2, 0.2]$ interval. This, as described in Primiceri (2005), ensures the convergence of the algorithm, its satisfactory performance and, therefore, its suitability to conduct econometric inference.

A.5 Responses to Non Labour Supply Shocks

For completeness, this appendix aims at showing the impulse responses of the variables in the baseline model to a business cycle, a wage bargaining and a residual shock. The shocks are all normalised to increase the real gross domestic product by 1% in median on impact.

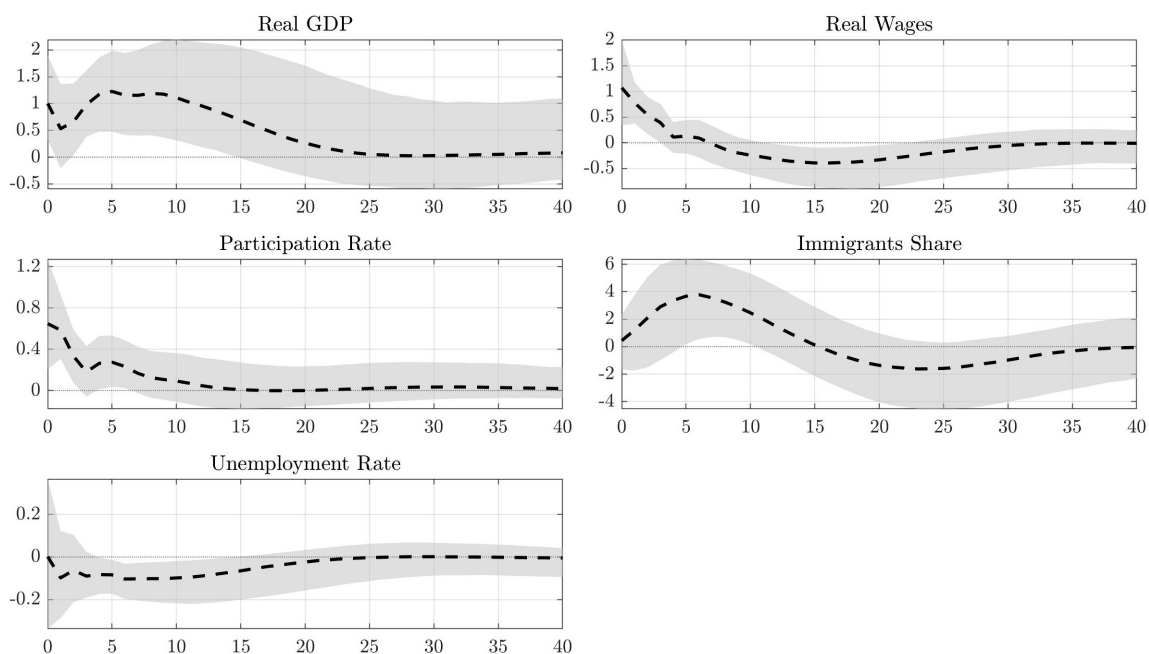


Figure 11: Impulse Responses to a Business Cycle Shock at each Horizon in the Baseline Model - Posterior Median Values (Black Dashed Lines) and 68% Bands (Grey Shaded Areas)

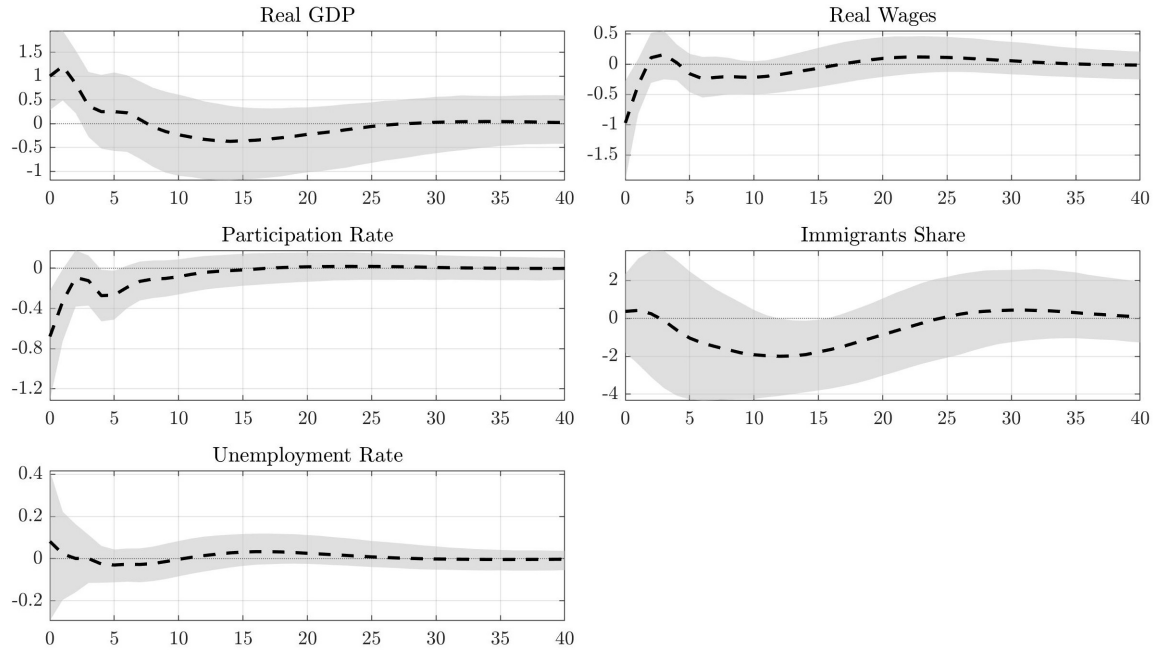


Figure 12: Impulse Responses to a Wage Bargaining Shock at each Horizon in the Baseline Model - Posterior Median Values (Black Dashed Lines) and 68% Bands (Grey Shaded Areas)

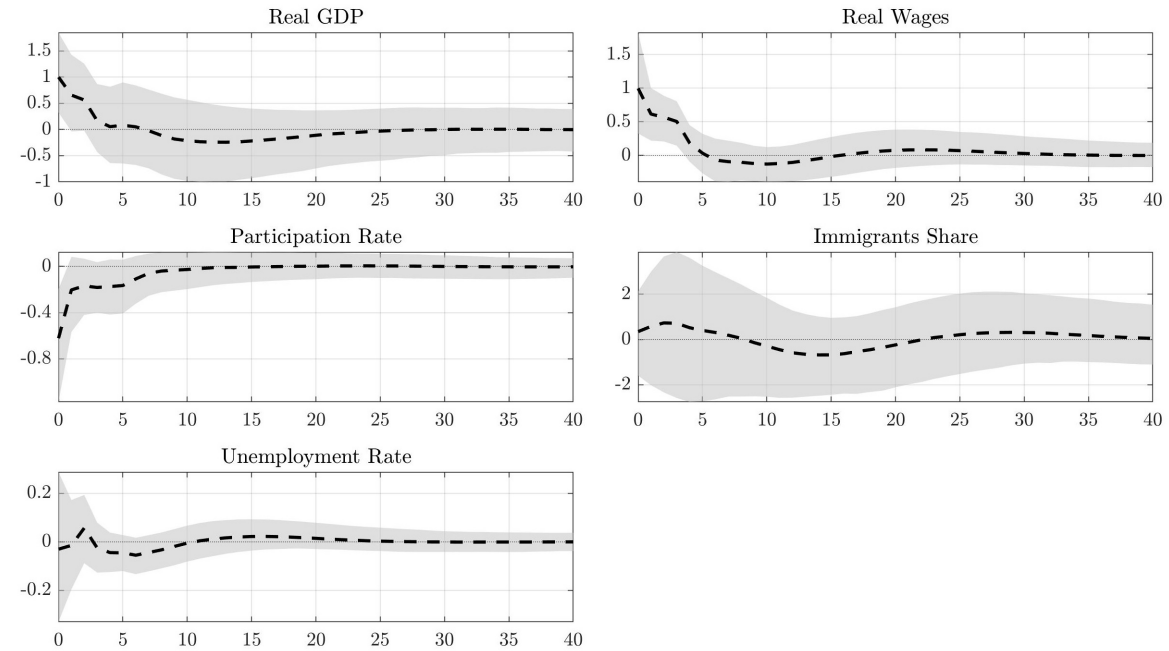


Figure 13: Impulse Responses to a Residual Shock at each Horizon in the Baseline Model - Posterior Median Values (Black Dashed Lines) and 68% Bands (Grey Shaded Areas)