Erasmus University Rotterdam
Erasmus School of Economics
Bachelor Thesis Econometrics and Operations Research

Mixed-Frequency GDP Nowcasting with Machine Learning Regularised PMIDAS regressions

Rochelle Busman (492443)

Abstract: This thesis enhances GDP nowcasting by integrating regularisation methods Ridge, LASSO, and Elastic Net into the panel mixed data sampling (PMIDAS) framework of (Fosten & Nandi, 2023) with common correlated effects for mixed-frequency data. Using OECD Main Economic Indicators database with business survey predictors, we nowcast real GDP growth of quarterly and yearly frequency for a large set of countries. We compare this to the root mean square forecast error of traditional Ordinary Least Squares and autoregressive estimation. Results show that pooled models outperform non-pooled models and that regularised models have higher prediction accuracy, with Elastic Net (weight 0.8) and LASSO excelling in yearly GDP nowcasts and Elastic Net (weight 0.2) and Ridge in quarterly nowcasts. We find that countries Israel and Türkiye have the highest distribution of nowcast errors, and that increasing the size of shrinkage parameter lambda lowers the number of selected variables for pooled Lasso. In addition, further analysis is conducted on the importance of predictor variables by means of selection frequency and relative size of the coefficients. We recommend to use Elastic Net for future nowcasting as this model combines the strengths of both Ridge and Lasso, while effectively addressing their limitations. This thesis provides valuable insights for econometric modelling and policy making by reducing overfitting and improving processing of high-dimensional mixed-frequency panel data.

Supervisor: Dr. A. Pick
Second assessor: PA D. Opschoor
Date final version: June 28, 2024

1 Introduction

This thesis aims to enhance nowcasting of real Gross Domestic Product (GDP) growth rates by integrating machine learning (ML) regularisation methods for shrinkage and lag selection. Nowcasting, predicting the present, very recent past and close future, often deals with ragged edge data due to the differences in publication lags of many economic predictors between countries. We extend the panel mixed data sampling approach (PMIDAS) introduced by Fosten and Nandi (2023) by applying penalised regression methods Ridge, LASSO, and Elastic Net to data with a ragged edge and cross-sectional dependence. In the analysis of quarter-on-quarter (Q-O-Q) and year-on-year (Y-O-Y) GDP growth, using a panel data set encompassing a large set of countries featuring six monthly key predictors, shrinkage estimators are used to generate and improve nowcasts. The goal is to achieve superior out-of-sample predictions compared to traditional Ordinary Least Squares (OLS) and time series autoregressive (AR1) estimation, which is measured by the root mean squared forecast error (RMSFE).

The importance of real-time data analysis for predicting key economic indicators like gross domestic product is common knowledge. Where many studies have traditional single-country nowcasting models, Baltagi (2008) and Wang, Zhang and Paap (2019) find evidence that there are benefits from pooling slope coefficients of countries by means of panel data, using pooled panel forecasts. Foroni, Marcellino and Schumacher (2015) find that if the difference in frequency is small enough, which is the case with monthly to quarterly and quarterly to yearly frequencies, distributed lag functions are not necessary, leading to the use of the UMIDAS model.

Build upon these theories, Fosten and Nandi (2023) develop the panel data extension of the mixed-frequency vector autoregressive model from Ghysels (2016), resulting in the existing PMIDAS framework. This framework has been further adjusted with a lagged modification of the common correlated effects (CCE) approach of Chudik and Pesaran (2015). Siliverstovs (2016) and Babii, Ball, Ghysels and Striaukas (2020) find that using machine learning to perform shrinkage and lag selection improves the model and accuracy of predictions, which motivates us to further expand estimation of the current PMIDAS model. The aim is to reduce overfitting and model complexity with regularised regression methods Ridge, Lasso, and Elastic Net. Penalised regression methods are beneficial when there is a high correlation between predictors and a significant risk of overfitting the data, which can result in poor prediction performance. Limiting overfitting of the training data by penalising parameter estimates allows us to reduce model complexity, possibly improving the out-of-sample predictions. Previous studies have applied machine learning methods to nowcasting of GDP growth(Zhang, Ni & Hao, 2023; Richardson, Mulder & Vehbi, 2018; Kant, Pick & Winter, 2022). Where Beyhum and Striaukas (2023) introduce Ridge-MIDAS within sparse regression methods, Liu, Liu, Li and Wen (2021) nowcast tourism demand by a penalty-based LASSO-MIDAS model. Tiffin (2016) focuses on Elastic Net regression for nowcasting.

This research is important for both scientific and practical purposes. It improves theory in the field of econometrics by comparing machine learning methods with more traditional estimation methods to make nowcasting more accurate. This can lead to models that handle high-dimensional data and multicollinearity more effectively. In general, Ridge regression reduces overfitting and Lasso reduces model complexity by means of feature selection. Improved prediction methods for panel data are especially valuable for policymakers, central banks, and economic analysts, who need precise and timely forecasts for key economic indicators like GDP growth and inflation. Accurate nowcasts of these indicators can greatly improve decision-making processes for monetary policy, fiscal planning, and economic stability.

It is necessary to research this, as current nowcasting models often struggle with issues such as overfitting and multicollinearity, leading to lower prediction accuracy. Economic indicators such as GDP are influenced by many variables, which makes it hard for models to distinguish actual patterns in the data set instead of noise. In addition, GDP growth and predictors for countries are often influenced by common regional or global economic factors, which may to high correlation of the dependent and independent variables.

Traditional estimation methods like OLS may not handle the complexity and interdependence of large panel data sets well, and by using penalised regression methods like Ridge, LASSO, and Elastic Net, complexity can be reduced and interpretation simplified. Current available information on this subject incorporating mixed-frequency data and cross-dependence is limited, but does exist. A thorough review of existing literature did not reveal any previous studies applying penalised regression methods to mixed-frequency Panel MIDAS frameworks with heterogeneous coefficients, making this research innovative.

The aim of this thesis is to research how integrating regularisation for shrinkage and lag selection with penalised regression methods Ridge, LASSO and Elastic Net to data with a ragged edge and possible cross-sectional dependence enhances nowcasting accuracy. This is analysed by means of comparing out-of-sample prediction accuracy to the PMIDAS model by Fosten and Nandi (2023) and time series AR measured with root mean squared forecast error (RMSFE). In addition, further analysis is conducted on pooling the ML-adjusted PMIDAS model and comparing the results to those produced for individual countries separately. Furthermore, deeper analysis of the penalisation weights for Elastic Net, individual country distributions of the nowcast errors, and the size of the shrinkage parameter lambda is performed. In addition to that, feature selection by Lasso and coefficient size are analysed for both target variables.

The first main finding is that pooled models consistently exhibit lower RMSFE compared to non-pooled models, and that data frequency has a high impact on the performance of the models, where models predicting quarterly GDP growth demonstrate significantly lower RMSFE than those predicting yearly GDP growth. Secondly, for yearly GDP growth, pooled Lasso and Elastic Net (= 0.8) show superior predictive accuracy, and pooled Ridge and Elastic Net (= 0.2) outperform other methods for quarterly GDP growth. The third main finding is that Israel and Türkiye present the largest distribution of nowcast errors and that a significant relationship exists between the shrinkage parameter and the number of selected variables for Pooled Lasso. This indicates that strong shrinkage can lead to underfitting and lower predictive accuracy for Q-O-Q GDP growth. Finally, we recommend to use Elastic Net for new data sets because of its flexibility in combining the penalties of both Ridge and Lasso. The predictive accuracy of Elastic Net is comparable to the other models, while effectively addressing the limitations of both Ridge and Lasso.

This thesis is further organised as follows. Section 2 discusses the methodology of the benchmark PMIDAS model and penalised regression models. Section 3 provides a detailed overview of the data set used in this study, including a deeper analysis of the predictors and target variables. The results of the study are presented in Section 4. This section starts with an analysis of the nowcast performance for different alpha values in the Elastic Net model to select the best model for further comparison. Following this, this section compares the nowcast performance of all estimation methods and conducts a deeper analysis on nowcast errors, shrinkage parameter lambda, individual country distributions and relative importance of the predictors. We examine the robustness of these results is by including additional predictors and using different out-of-sample split ratios. Finally, section 5 provides a discussion of the research, its implications, and concludes the thesis by summarising the key findings and suggesting directions for future research.

2 Methodology

2.1 Introduction to PMIDAS model

The PMIDAS nowcasting model introduced by Fosten and Nandi (2023) allows for heterogeneity and common correlated effects (CCE). We apply this model to mixed-frequency data with two ratios for high to low frequency. The first ratio is monthly-to-quarterly with ratio $q = 3$ and the second ratio is quarterly-to-yearly with $q = 4$. In this section, we will focus on the first ratio for a frequency mix of monthly-to-quarterly. However, the equations can easily be adjusted to the second frequency mix. Equation 1 specifies for every country i the vector $X_{i,t}^M$ that stacks all three monthly observations $x_{i,t}^M$ of the predictor variable for every quarter t. This theory by Ghysels (2016) allows us to combine data from different frequencies.

$$X_{i,t}^M = \begin{pmatrix} x_{i,t}^M \\ x_{i,t-\frac{1}{3}}^M \\ x_{i,t-\frac{2}{3}}^M \end{pmatrix}, \tag{1}$$

As specified by Fosten and Nandi (2023) we define variable of interest quarterly real GDP growth $y_{i,t}$, where $i = 1, 2, ..., N$ is set of cross-sectional units, specifically countries. Time is defined by $t = 1, 2, ..., T$, where T is the total number of quarters in the data set. This approach allows addressing the ragged edge by utilising the most current data available, as the latest available observations may vary per country through the staggered release of new information. Variable v denotes the specific day within the nowcasting period for which the nowcast is being made. Variable d_{iv} represents the latest available quarterly lag of the target variable $y_{i,t}$ for country i as of day v within the nowcasting quarter. This variable captures the most recent quarter's data that influences the nowcast. In addition, m_{iv} indicates the latest available monthly lag for the predictor variable $X_{i,t}^M$ for country i on day v of the nowcasting quarter. When $m_{iv} = 0$, it signifies that all three months of data for the quarter are available for use in the prediction.

The final equation used for nowcasting by Fosten and Nandi (2023) is specified in Equation (2), where c_{vi} are the country-specific fixed effects, ϕ_{vi} is the autoregressive lag coefficient, β_{vi}' comprises of individual-specific slope coefficients for the predictor variable $X_{i,t-\frac{m_{i,v}}{3}}^M$ in Equation (1), and γ_{vi} is the coefficient for the cross-sectional average of predictor variable $\overline{X}_{i,t-\frac{m_{i,v}}{3}}^M$ and its lags. The sum is taken from lag 0 to lag truncation $p_T = T^{1/3}$, defined by Chudik and Pesaran (2015) over δ'_{vil} infinite lag polynomial multiplied by the cross-sectional weighed average of $\overline{z}^M_{t-l,v}$. Finally, the error term is defined by $e_{v,i,t}$. The first benchmark model for comparison is estimation of Equation (2) with Ordinary Least Squares.

$$y_{i,t} = c^*_{vi} + \phi_{vi} y_{i,t-d_{iv}} + \beta'_{vi} X^M_{i,t-\frac{m_{i,v}}{3}} + \sum_{l=0}^{P_T} \delta'_{vil} \overline{z}^M_{t-l,v} + \gamma'_{vi} \overline{X}^M_{i,t-\frac{m_{i,v}}{3}} + e_{v,i,t} \tag{2}$$

To provide a second benchmark for comparison, we estimate a simple AR(1) model for each country i with Equation (3).

$$y_{i,t} = \alpha_i + \rho y_{i,t-d_{iv}} + \epsilon_{i,t} \tag{3}$$

2.2 Penalised regression methods

Equation (2) is estimated with ordinary least squares in the paper by Fosten and Nandi (2023). In this thesis, we will apply penalised regression for parameter estimation to this equation through regularisation methods Ridge, LASSO and Elastic Net to reduce overfitting and improve predictive performance. With these methods, a penalty term is introduced, which is a function of the coefficients, multiplied by regularisation parameter $\lambda \geq 0$. This value of λ is also called shrinkage parameter, and manages the amount of shrinkage. A higher value of λ corresponds to more shrinkage and results in smaller coefficients. Regularisation methods introduce the bias-variance trade-off. Minimising the residual sum of squares while placing constraints on the parameters introduces a bias. This constraint on the coefficient limits the likelihood space, which makes the probability that the coefficient is equal to true value smaller (Blei, 2015). At the same time, variance is reduced because the coefficient can take on fewer possible values. Decreasing λ will make the penalty term weaker, resulting in lower bias and higher variance. Shrinkage method Lasso conducts subset selection of variables and lags with a L1 penalty, which may lead to lower prediction errors in nowcasting (Hastie, Tibshirani & Friedman, 2013). Ridge regression imposes a L2 penalty on the size of the coefficients by minimising the residual sum of squares. It may shrink the coefficients until they are very close to zero, but Ridge will make them equal to 0, indicating that they will not remove any variables. The objective function of the Ridge regression is shown in Equation (4), where penalty term L2 is equal to the squared sum of coefficients, including coefficients β_{vi}, ϕ_{vi}, δ_{vil}, and γ_{vi}.

$$\begin{aligned}\min_{\beta,\phi,\delta,\gamma} \Bigg\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(y_{i,t} - c^*_{vi} - \phi_{vi} y_{i,t-d_{iv}} - \beta'_{vi} X^M_{i,t-\frac{m_{i,v}}{3}} - \sum_{l=0}^{P_T} \delta'_{vil} \overline{z}^M_{t-l,v} - \gamma'_{vi} \overline{X}^M_{i,t-\frac{m_{i,v}}{3}} \right)^2 \\ + \lambda_1 \left(|\beta_{vi}|^2 + |\phi_{vi}|^2 + |\gamma_{vi}|^2 + \sum_{l=0}^{P_T} |\delta_{vil}|^2 \right) \Bigg\}\end{aligned} \tag{4}$$

With Lasso, the L1 penalty forces some or many of the coefficient estimates to zero, if it thinks that they are not relevant enough. This means that Lasso can replace other methods that select only a few important variables, which makes the model simpler and easier to interpret.

The L1 Lasso penalty results in nonlinear solutions of the target variable, which means that a closed-form solution does not exist(Hastie et al., 2013). Equation (5) shows the objective function with LASSO, where the penalty term is equal to the absolute sum of the coefficients, including the coefficients β_{vi}, ϕ_{vi}, δ_{vil}, and γ_{vi}. This method can select variables and lags of the predictor variables by shrinking their coefficients to zero.

$$\min_{\beta,\phi,\delta,\gamma} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(y_{i,t} - c^*_{vi} - \phi_{vi} y_{i,t-d_{iv}} - \beta'_{vi} X^M_{i,t-\frac{m_{i,v}}{3}} - \sum_{l=0}^{P_T} \delta'_{vil} \overline{z}^M_{t-l,v} - \gamma'_{vi} \overline{X}^M_{i,t-\frac{m_{i,v}}{3}} \right)^2 + \lambda_2 \left(|\beta_{vi}| + |\phi_{vi}| + |\gamma_{vi}| + \sum_{l=0}^{P_T} |\delta_{vil}| \right) \right\} \quad (5)$$

Elastic Net combines the strengths of Ridge and LASSO regularisation by applying both forms of penalty terms L1 and L2, shown in Equation (6). We select hyperparameter α by means of predictive performance measured with RMSFE. This value indicates how much weight is given to both penalties (Hastie et al., 2013). Alpha is a value between 0 and 1. If alpha is equal to 0, Elastic Net transforms into Ridge Regression, and if alpha is equal to 1, Elastic Net becomes Lasso.

$$\min_{\beta,\phi,\delta,\gamma} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(y_{i,t} - c^*_{vi} - \phi_{vi} y_{i,t-d_{iv}} - \beta'_{vi} X^M_{i,t-\frac{m_{i,v}}{3}} - \sum_{l=0}^{P_T} \delta'_{vil} \overline{z}^M_{t-l,v} - \gamma'_{vi} \overline{X}^M_{i,t-\frac{m_{i,v}}{3}} \right)^2 + \lambda \left[\alpha \left(|\beta_{vi}| + |\phi_{vi}| + |\gamma_{vi}| + \sum_{l=0}^{P_T} |\delta_{vil}| \right) + (1-\alpha) \left(|\beta_{vi}|^2 + |\phi_{vi}|^2 + |\gamma_{vi}|^2 + \sum_{l=0}^{P_T} |\delta_{vil}|^2 \right) \right] \right\} \quad (6)$$

After obtaining parameter estimates for β, ϕ, δ, and γ, the nowcast on day v of quarter T for every country i can be estimated by means of the conditional mean of $\hat{y}_{i,T,v}$, shown in Equation (7).

$$\hat{y}_{i,T,v} = \hat{c}^*_{vi} + \hat{\phi}_{vi} y_{i,T-d_{iv}} + \hat{\beta'}_{vi} X^M_{i,T-\frac{m_{i,v}}{3}} + \sum_{l=0}^{P_T} \hat{\delta}_{vil} \overline{z}^M_{T-l,v} + \hat{\gamma}_{vi} \overline{X}^M_{i,T-\frac{m_{i,v}}{3}} \quad (7)$$

The optimal value regularisation parameter λ for all penalised regression models is chosen by means of 10-fold cross-validation. If this parameter is chosen to be equal to 0, the ML-adjusted forecast equations correspond to Equation (7). For all estimation methods, we apply two types of assumptions to both target variables. Assuming full heterogeneity of coefficients with non-pooled models, each equation is estimated separately for each country. Fosten and Nandi (2023) find that these models under perform in predictive performance when compared to pooled models, possibly due to overfitting. Pooled models impose homogeneity on all slope coefficients across countries, but still allow for individual-specific fixed effects by means of intercept.

2.3 Evaluation of model accuracy

Finally, the accuracy of the model is measured by means of root mean squared forecast error (RMSFE) defined in Equation (8). The nowcast error $\hat{e}_{i,t}$ is equal to the difference between real values $y_{i,t}$ and the predicted nowcasts.

$$\text{RMSFE}_{v,i} = \sqrt{\frac{1}{P}\sum_{t=R}^{T}\hat{e}^2_{v,i,t}} \tag{8}$$

3 Data

As the aim of this research is to improve the forecast accuracy of the the PMIDAS model from Fosten and Nandi (2023), the model estimated with Ordinary Least Squares is taken as main benchmark and we use the same data set to be able to make an honest comparison. This data set stems from the Organisation for Economic Co-operation and Development (OECD), that publishes the Main Economic Indicators (MEI) database for key economic variables (OECD, 2024). The main target variables for nowcasting are the real GDP growth rates with two frequencies: year-on-year (Y-O-Y) and quarter-on-quarter (Q-O-Q). Different from prior research, nowcasts are made for individual countries separately, and they are not aggregated to a global level of real GDP growth. This is possible because the predictor variables are available for all countries individually. Table 5 in the Appendix lists the predictor variables, available time period, sample size and the specific countries for which they are available. The monthly predictors are the industrial production index (IP) and business survey indices for manufacturing (BSM), services (BSS), consumer (CON), trade (TRA) and construction (CST). New predictor variables are created by taking the average and including up to four lags. The data set contains a large set of emerging and advanced countries, where for some countries the predictors are not published for the same time period or not published at all. One of the findings by Fosten and Nandi (2023) is that it is better to choose a single predictor variable than a combination of predictors in the PMIDAS model, where the BSM variable as the only external predictor yields the best results. Adding more information by allowing for multiple predictor variables yields a lower prediction accuracy, which can possibly be explained by the need to estimate a higher number of coefficients and the high correlation between the predictors.

Figure 1 shows the real GDP growth rates of all countries over time for both target frequencies year-on-year and quarter-on-quarter over time. Most countries show a steady growth rate with small fluctuations between -0.1 and 0.1. During the global financial crisis in 2008 and the COVID-29 pandemic in 2020, a downward trend is depicted for all countries. Noteworthy are the real GDP growth rates for Israel, which has significantly higher values in 2015, with a maximum value of 2.5%. This can be explained by their strong growth of private consumption, resulting from high employment, low interest rates and low import prices (International Monetary Fund, 2015). Figure 1a for year-on-year GDP growth rates depicts higher volatility than figure 1b, also demonstrated by their cross-country standard deviations of 0.034 and 0.013, respectively. In general, year-on-year growth trends tend to be more volatile than quarter-on-quarter growth rates (ECB, 2000). However, this seems not to be the case in figure 1, possibly because they are influenced by more events over a longer period of time. In addition, year-on-year GDP growth rates also capture the cumulative developments of the quarter-on-quarter growth rates (ECB, 2000).

Figure 1: Real GDP growth over time for targets year-on-year and quarter-on-quarter

(a) Real GDP growth for Y-O-Y over time **(b)** Real GDP growth for Q-O-Q over time

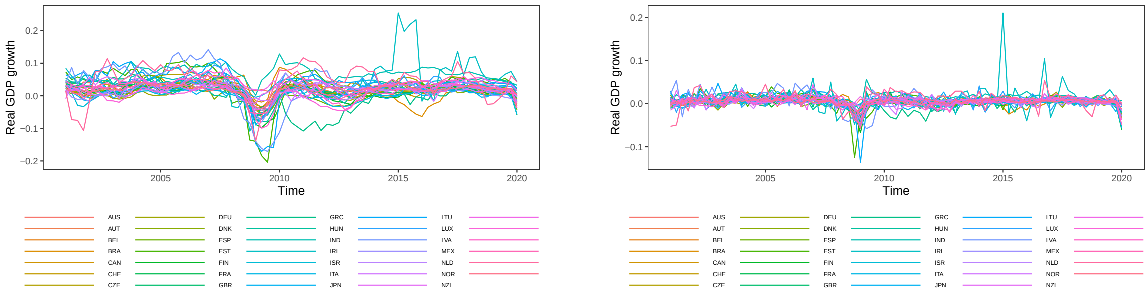

Note: Real GDP growth rates for individual countries from Q1 2001 to Q4 2020.

Table 6 in the Appendix shows the cross-correlation for GDP and all six predictor variables for yearly observations of the common period from 2003 to 2005. Target variable GDP has low correlation with all predictor variables. However, this does not appear to be an issue as they can still strongly predict GDP growth, shown by Fosten and Nandi (2023). Strong correlation between business surveys Manufacturing and Services is demonstrated by a high correlation value of 0.955. There is a moderate positive correlation of 0.452 between Construction and Trade, and of 0.334 between the Industrial Production Index and Manufacturing. This risk of multicollinearity creates a higher relevance of applying penalised regression methods to reduce its effect and improve prediction accuracy. In addition, low correlation between the other predictors suggests that they capture different aspects of economic trends, possibly leading to a well-performing model with multiple predictors. At the same time, it is important to reduce model complexity by reducing the dimensionality of the model to avoid overfitting and enhance out-of-sample predictions. This will be performed by regularisation methods Ridge, Lasso and Elastic Net. Table 7 in the Appendix shows the cross-correlation between all predictors derived from GDP and Business Manufacturing Services, including their respective averages and lags. Many pairs of predictors exhibit high correlation values, which can inflate the variance of coefficient estimates and reduce predictive power of the models. In addition, extremely high VIF values are calculated for this predictor matrix, which confirms that multicollinearity is present. This validates our approach of applying regularisation methods to mitigate these effects.

Figure 17 in the Appendix displays the cross-country distributions of GDP and three main economic indicators Manufacturing, Services and Industrial production. It can be seen that the economic cycles of the variables move approximately the same, with notable dips during the global financial crisis years. Despite some variations, the common time patterns in the figure demonstrate their suitability as GDP growth predictors.

Figure 2 demonstrates the autocorrelation functions of GDP and the main predictor variables. All show significant autocorrelation at several lags, which indicates that past values are strongly correlated with future values. Where GDP has significant autocorrelation up to four lags, autocorrelation for Manufacturing, Services and Industrial production ranges between 9 to

16 lags. The long presence of significant autocorrelation and slow decline imply that the series are not stationary.

Machine learning regularisation techniques tend to be useful as they can select the most relevant lags by shrinking the coefficients of less important lags to zero. This helps to better capture underlying patterns and improve the nowcast quality. This nowcast quality is evaluated by means of a pseudo out-of-sample experiment with recursive estimation of all five methods. The sample is split into a training set R and test set P, with full sample $T = R + P$. To evaluate robustness of the results, the evaluation window is set to 3 out-of-sample split ratios, namely 0.2, 0.3, and 0.4, each corresponding to 20%, 30%, 40% of the total sample size used for evaluation. We adopt the approach from Fosten and Nandi (2023) where nowcasts are generated for each day over a 155-day window from the beginning of each quarter, including backcasting from day 91 onward. The pseudo-calendar replicates average release days of all variables with assumptions for 30-day months and 90-day quarters. All results presented in this thesis are generated using R, utilising the *glmnet* package for implementation of the penalised regression methods. Shrinkage parameter lambda is estimated by means of cross-validation with *lambda.1se*, ensuring that the mean squared error is within one standard error of the minimum (Hastie, Qian & Tay, 2023).

Figure 2: Autocorrelation Functions of GDP and three main predictor variables

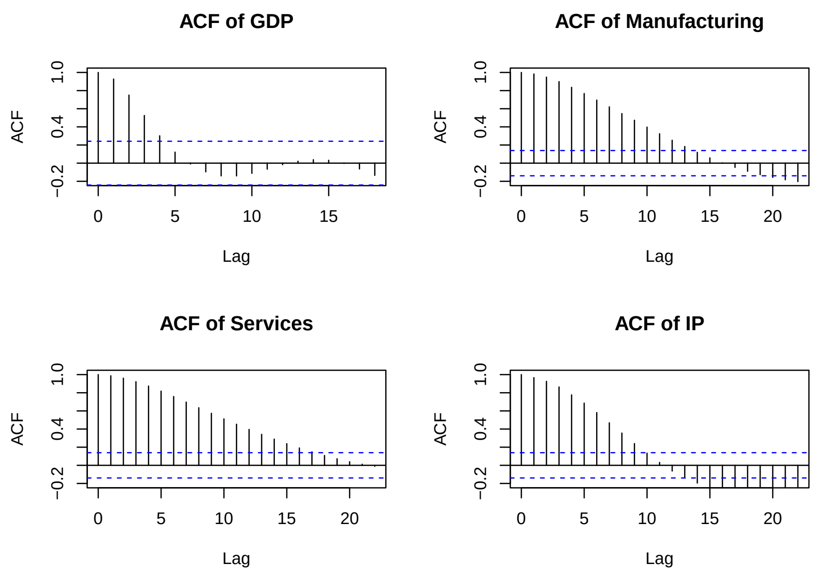

Note: Figure displays the autocorrelation functions for GDP growth and three predictor variables.

4 Results

4.1 Nowcast performance alpha parameters Elastic Net

The Elastic Net model is a combination of Ridge regression and Lasso and is estimated with four different values for the parameter alpha, namely 0.2, 0.4, 0.6, and 0.8. When alpha is equal to 0.0, Elastic Net is equal to Ridge regression with L2 penalty, and when alpha is equal to 1, Elastic Net reduces to the Lasso estimation method with the L1 penalty. Based on figure 3, we select alpha with the best out-of-sample prediction performance for target variables year-on-year and quarter-on-quarter GDP separately. The first result that we can conclude from this figure is that pooled models have better predictive accuracy than non-pooled models. We will further analyse this result in section 4.2. Figure 3a shows that for Y-O-Y the nowcast RMSFE of Elastic Net is the lowest when alpha is equal to 0.8. This is based on the blue solid line for pooled Elastic Net of alpha 0.8 that plots the RMSFE, which is clearly lower than the other models after day 45. Figure 3b demonstrates that for Q-O-Q the nowcast RMSFE of Elastic Net is the lowest when alpha is equal to 0.2, which is plotted by the solid red line for Pooled Elastic Net. From day 0, this estimation method has a lower RMSFE than the other models. These results are generated for the out-of-sample split of 30%. To ensure the robustness of results with regards to the out-of-sample split ratio, more figures are displayed for different sample splits in the Appendix in section B. Figure 20 shows the nowcast RMSFE for ratios 0.2 and 0.4 for year-on-year GDP growth, and figure 21 for quarter-on-quarter GDP growth. When comparing the level of RMSFE on day 0 for the different out-of-sample split ratios figures, it can be concluded that when the out-of-sample split is lower, the average RMSFE over the nowcast period is lower. This can be explained by the fact that lower out-of-sample split ratio allows for better model training and captured variability in the test set. Model comparison in the next sections are performed by means of the best performing alpha for Elastic Net, where $\alpha = 0.8$ for Y-O-Y and $\alpha = 0.2$ for Q-O-Q GDP growth.

Figure 3: Elastic Net nowcast RMSFE of alpha parameters for year-on-year and quarter-on-quarter GDP

(a) Target Y-O-Y

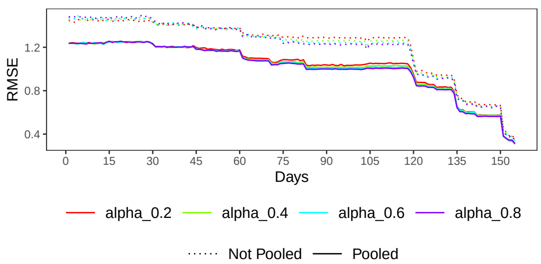

(b) Target Q-O-Q

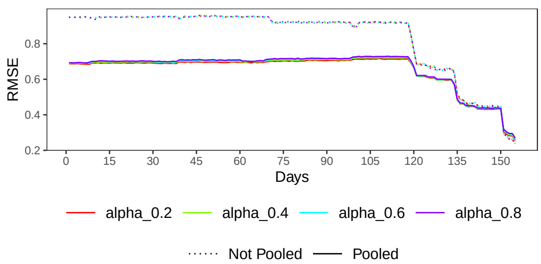

Note: This figure plots for each alpha ($\alpha = 0.2, 0.4, 0.6, 0.8$) of Elastic Net the cross-country average root mean squared forecast error (RMSFE) over the full nowcast period of 155 days. The predictor variable is Business Survey Manufacturing, and the target variables are year-on-year and quarter-on-quarter GDP, with out-of-sample split = 30%.

4.2 Nowcast performance comparison

After selecting the best performing alphas for Elastic Net, we will compare the prediction accuracy of all five estimation methods OLS, Ridge, Lasso, Elastic Net and time series AR, using only predictor variable business surveys Manufacturing to predict GDP growth. Table 1 and table 2 report the cross-country quantiles of RMSFE for year-on-year GDP growth, where table 1 is for non-pooled models and 2 is for pooled models. If the models are pooled, this means that homogeneity is imposed on the slope coefficients for every country individually, while still allowing for heterogeneity with individual-specific-effects. Model time series AR is reported in both tables to be able to make a full comparison. It can be concluded from both tables 1 and 2 that penalised regression methods Ridge, Lasso and Elastic Net show consistently lower

RMSFE across all quantiles compared to OLS and AR. Both Lasso and Elastic Net with alpha 0.8 perform well, and the values are very close to each other, but Elastic Net outperforms Lasso with a generally lower RMSFE. However, with a weight of 0.8 for the L1 parameter, Elastic Net is nearly equivalent to Lasso, making the models almost identical. When comparing table 1 to table 2, it is clear that pooled models strictly outperform non-pooled models with lower RMSFE, shown by a difference of at least 0.2 on day 1 for every model.

Table 1: Y-O-Y GDP nowcast RMSFE by quantile for all non-pooled models

	Methods														
	OLS			**Ridge**			**Lasso**			**Elastic Net**			**TS AR**		
Days	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**
1	1.079	1.366	1.688	0.959	1.169	1.510	0.928	1.190	1.583	0.925	1.159	1.551	0.987	1.131	1.464
16	1.079	1.366	1.688	0.968	1.173	1.496	0.913	1.151	1.543	0.934	1.183	1.610	0.987	1.131	1.464
31	1.066	1.353	1.568	0.916	1.175	1.412	0.916	1.160	1.392	0.903	1.158	1.400	0.908	1.064	1.302
46	1.009	1.261	1.493	0.893	1.102	1.409	0.908	1.115	1.354	0.890	1.091	1.348	0.887	1.049	1.287
61	0.929	1.157	1.480	0.857	1.083	1.380	0.828	1.009	1.260	0.829	1.032	1.288	0.812	0.925	1.287
76	0.746	0.962	1.278	0.936	1.078	1.314	0.803	0.943	1.194	0.806	0.948	1.206	0.763	0.908	1.255
91	0.738	1.025	1.266	0.836	1.079	1.324	0.809	0.954	1.210	0.793	0.946	1.196	0.733	0.895	1.255
106	0.738	1.025	1.266	0.844	1.062	1.318	0.804	0.938	1.206	0.797	0.950	1.185	0.733	0.895	1.255
121	0.460	0.705	1.061	0.603	0.943	1.166	0.519	0.805	1.117	0.527	0.802	1.132	0.475	0.763	1.049
136	0.000	0.000	1.047	0.000	0.000	1.039	0.000	0.000	0.925	0.000	0.000	1.040	0.000	0.000	0.832
151	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: This table presents the Root Mean Square Forecast Error (RMSFE) of GDP nowcasts by quantile (25%, 50%, 75%) across models, including OLS, Ridge, Lasso, Elastic Net, and TS AR, over all days in the nowcasting period. The out-of-sample split ratio is 30%.

Table 2: Y-O-Y GDP nowcast RMSFE by quantile for all pooled models

	Methods														
	OLS			**Ridge**			**Lasso**			**Elastic Net**			**TS AR**		
Days	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**
1	0.785	1.057	1.313	0.699	0.985	1.338	0.695	1.003	1.231	0.685	0.998	1.236	0.987	1.131	1.464
16	0.796	1.044	1.337	0.710	0.995	1.334	0.713	1.059	1.241	0.708	1.035	1.245	0.987	1.131	1.464
31	0.727	1.012	1.365	0.711	0.947	1.313	0.649	0.932	1.199	0.653	0.939	1.215	0.908	1.064	1.302
46	0.657	0.913	1.183	0.672	0.923	1.235	0.611	0.911	1.151	0.618	0.914	1.143	0.887	1.049	1.287
61	0.615	0.815	1.077	0.631	0.865	1.136	0.572	0.806	1.053	0.572	0.808	1.062	0.812	0.925	1.287
76	0.560	0.743	1.027	0.590	0.811	1.111	0.562	0.730	0.973	0.573	0.728	0.974	0.763	0.908	1.255
91	0.551	0.732	1.001	0.584	0.793	1.102	0.545	0.712	0.986	0.556	0.721	0.978	0.733	0.895	1.255
106	0.544	0.729	1.027	0.601	0.821	1.121	0.558	0.709	0.998	0.571	0.708	0.994	0.733	0.895	1.255
121	0.404	0.625	0.995	0.404	0.774	1.088	0.443	0.644	0.978	0.446	0.652	0.978	0.475	0.763	1.049
136	0.000	0.000	0.799	0.000	0.000	0.912	0.000	0.000	0.704	0.000	0.000	0.721	0.000	0.000	0.832
151	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: Please see the footnote of Table 1.

Table 3 and 4 demonstrate the quantiles of RMSFE for target variable quarterly GDP growth (Q-O-Q). Comparing these tables to table 1 and 2, it can be seen that the RMSFE is significantly lower for quarterly GDP than for yearly GDP. This result can be attributed to various factors. Year-on-year GDP growth exhibits higher volatility compared to quarter-on-quarter GDP growth, as shown in section 3. This makes it more difficult for all estimation methods to make accurate predictions, which may lead to higher RMSFE values. In addition, predictor variable Business Survey Manufacturing may have more relevance and better fit for short term than long term growth predictions. Longer time periods between observations can mask short-term differences, making it more difficult for the estimation methods to detect and respond to changes quickly. When comparing table 3 and 4 to each other, we come to the same conclusion that pooled models perform better overall, as table 4 has lower RMSFE values for all three quantiles on every nowcast day. The penalised regression methods outperform OLS and TS in both tables. For quarterly GDP growth, it can be concluded that the best performing models are Ridge regression and Elastic Net with a value for alpha of 0.2. The quantiles of both methods are very close to each other for all nowcast days, however, Elastic Net shows superior performance due to the lowest RMSFE overall.

Table 3: Q-O-Q GDP nowcast RMSFE by quantile for all non-pooled models

	Methods														
	OLS			**Ridge**			**Lasso**			**Elastic Net**			**TS AR**		
Days	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**
1	0.713	0.898	1.254	0.589	0.738	1.026	0.588	0.731	1.025	0.588	0.738	1.026	0.595	0.736	1.032
16	0.713	0.898	1.254	0.588	0.738	1.026	0.586	0.731	1.026	0.589	0.738	1.026	0.595	0.736	1.032
31	0.732	0.890	1.247	0.588	0.737	1.026	0.588	0.738	1.026	0.588	0.733	1.012	0.595	0.732	1.032
46	0.682	0.886	1.247	0.585	0.738	1.026	0.585	0.738	1.041	0.585	0.738	1.015	0.588	0.732	0.982
61	0.720	0.931	1.246	0.585	0.731	1.010	0.585	0.731	0.984	0.585	0.731	0.985	0.588	0.739	0.982
76	0.647	0.721	1.124	0.585	0.718	0.978	0.585	0.728	0.962	0.585	0.723	0.965	0.588	0.742	0.982
91	0.633	0.748	1.085	0.585	0.726	0.965	0.585	0.715	0.999	0.585	0.725	0.949	0.588	0.742	0.982
106	0.633	0.748	1.085	0.585	0.729	0.957	0.585	0.730	0.978	0.585	0.716	0.942	0.588	0.742	0.982
121	0.322	0.643	0.806	0.365	0.580	0.859	0.365	0.585	0.795	0.365	0.585	0.834	0.380	0.614	0.805
136	0.000	0.000	0.731	0.000	0.000	0.717	0.000	0.000	0.621	0.000	0.000	0.689	0.000	0.000	0.742
151	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: Please see the footnote of Table 1.

Table 4: Q-O-Q GDP nowcast RMSFE by quantile for all pooled models

	Methods														
	OLS			**Ridge**			**Lasso**			**Elastic Net**			**TS AR**		
Days	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**	**25%**	**50%**	**75%**
1	0.431	0.545	0.701	0.330	0.487	0.609	0.361	0.506	0.643	0.354	0.492	0.642	0.595	0.736	1.032
16	0.439	0.543	0.722	0.330	0.478	0.607	0.389	0.520	0.633	0.369	0.501	0.637	0.595	0.736	1.032
31	0.446	0.545	0.728	0.328	0.484	0.599	0.384	0.507	0.637	0.364	0.497	0.633	0.595	0.732	1.032
46	0.397	0.549	0.667	0.331	0.479	0.597	0.404	0.524	0.620	0.382	0.503	0.618	0.588	0.732	0.982
61	0.401	0.546	0.662	0.325	0.481	0.587	0.390	0.521	0.606	0.376	0.504	0.600	0.588	0.739	0.982
76	0.382	0.542	0.673	0.343	0.486	0.594	0.400	0.536	0.621	0.384	0.511	0.602	0.588	0.742	0.982
91	0.422	0.578	0.700	0.344	0.487	0.593	0.428	0.533	0.621	0.411	0.523	0.613	0.588	0.742	0.982
106	0.424	0.571	0.701	0.351	0.496	0.585	0.418	0.558	0.653	0.401	0.529	0.616	0.588	0.742	0.982
121	0.317	0.514	0.667	0.260	0.484	0.578	0.283	0.495	0.582	0.276	0.504	0.593	0.380	0.614	0.805
136	0.000	0.000	0.573	0.000	0.000	0.495	0.000	0.000	0.509	0.000	0.000	0.503	0.000	0.000	0.742
151	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: Please see the footnote of Table 1.

The main prediction accuracy results in this thesis are shown in figure 4 where the nowcast average RMSFE for both target variables year-on-year and quarter-on-quarter GDP growth are demonstrated. These findings align with previous findings of the performance of all nine models, but are now shown graphically. Figure 4a shows that pooled Elastic Net with alpha 0.8 and Pooled Lasso have the best prediction accuracy for year-on-year GDP growth. Figure 4b shows that pooled Elastic Net with alpha 0.2 and pooled Ridge regression are best at predicting quarter-on-quarter GDP growth. Both are demonstrated by the lowest RMSFE values in general over the nowcast period of 155 days. Based on this, we can conclude that this difference in best performing models may be attributed to the difference in time comparison of GDP growth used in the nowcasts. The difference between the best performing models in both targets lies in the alpha of 1.0 for Lasso and 0.8 for Elastic Net (Y-O-Y) and 0.0 for Ridge and 0.2 for Elastic Net (Q-O-Q). These models are almost the same, except for the fact that Elastic Net employs both L1 and L2 regularisation, with the most weight given to either Ridge or Lasso. Lasso uses L1 regularisation, adding a penalty that is equal to the absolute value of the size of the coefficients, which has the effect of shrinking some coefficients to exactly zero. This elimination of lags or transformed variables of the predictor and target variable tends to be favourable when more volatile data is available, which is the case with year-on-year GDP. Yearly GDP encompasses a longer period and reflects long term changes, where the range of values is wider, and the number of significant events which affect GDP are higher. Shorter comparison periods such as quarter-on-quarter GDP reflect more short term changes. In addition, Q-O-Q growth rates are able to capture seasonal patterns in the data.

When estimation is performed by lasso, its ability to set coefficients to zero helps in focusing on the most important variables, by performing variable selection. This approach may have a better fit with Y-O-Y GDP growth, because less predictor variables may be relevant for this long term growth. Eliminating variables turns out to be the better approach. Ridge is not able to do that, as it may shrink coefficients towards zero, but never sets them equal to 0. Ridge regression uses L2 regularisation, which penalises the residual sum of squares of the coefficients. The result that Ridge regression performs better with quarter-on-quarter GDP may be explained by the fact that more transformed predictor variables are relevant for Q-O-Q GDP growth. Ridge shrinks, but retains all predictors, which turn out to be relevant for this target variable. Elastic Net combines both regularisation methods, and as the weights towards either L1 or L2 are large in the nowcasts, it is a logical consequence that its performance is approximately the same.

Figure 4: Nowcast average RMSFE of estimation methods for targets year-on-year and quarter-on-quarter GDP growth

(a) Y-O-Y nowcast average RMSFE **(b)** Q-O-Q nowcast average RMSFE

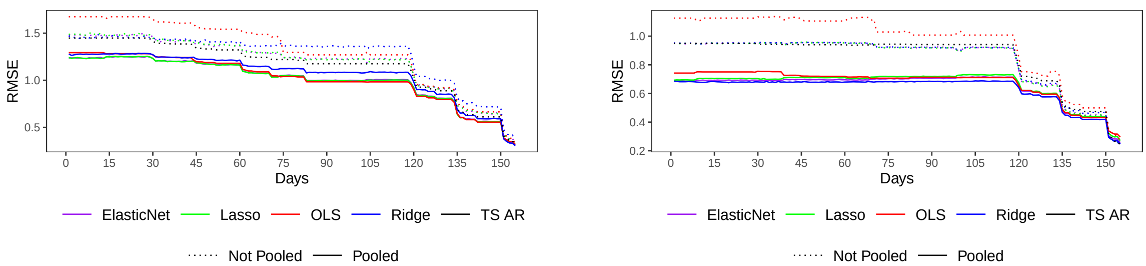

Note: Root mean squared forecast error (RMSFE) for full nowcast period. Out-of-sample split = 30%.

Important to mention is the bias-variance trade-off that all three regularisation methods introduce, which illustrates the balance between a model's capacity to minimise bias and variance. Increasing the model complexity, decreases the training error and the data is better fitted. However, if the data is fitted too well, this is called overfitting, leading to less generalisation and larger test error. Thus, increasing model complexity increases the variance and decreases bias (Hastie et al., 2013). To ensure robustness of results with regard to sample split, figures 18 and 19 in the Appendix show the results for out-of-sample ratios 0.2 and 0.4. These figures confirm previous findings about the best performing models.

Table 9 and 10 in the Appendix show the average RMSFE of the full nowcast period for all estimation methods and all three splits. Both tables confirm previous results that Lasso performs better for yearly GDP nowcasts and and Ridge for quarterly GDP nowcasts. Higher out-of-sample split ratios lead to higher average RMSFE. It can be concluded that pooled models outperform non-pooled models based on the order of the estimation methods and that time series AR is approximately in between. Pooled models assume that the coefficients are homogeneous across different countries. This is only helpful if there is not much heterogeneity and the differences between slopes of the countries are not large to begin with (Baltagi, 2008). Pooling countries simplifies the model and reduces variance of the estimated coefficients. In addition, instead of performing separate regressions for every country, pooling the countries increases the sample size in the regression, which leads to more precise estimates of the model parameters, reducing nowcast errors and resulting in lower RMSFE.

4.3 Distribution of nowcast errors across countries

Deeper analysis of the results is important, if we want to generalise the performance of these models to other data sets. This leads us to reviewing the nowcast error distributions for all countries in the panel individually. Figure 5 and 6 show the distributions of nowcast errors across countries for year-on-year and quarter-on-quarter, respectively. It can be concluded from these figures that the nowcast errors are relatively consistent across countries, with most nowcast errors around zero. Interesting to see is that Israel and Türkiye show a wider distribution of errors compared to all other countries with some very striking outliers, both for pooled and non-pooled models. This suggests that the economic dynamics of these countries are not well-captured by the current models, which may significantly impact the cross-country RMSFE discussed before.

Figure 5: Distribution of nowcast errors across countries for target year-on-year GDP growth

(a) Pooled models **(b)** Non-pooled models

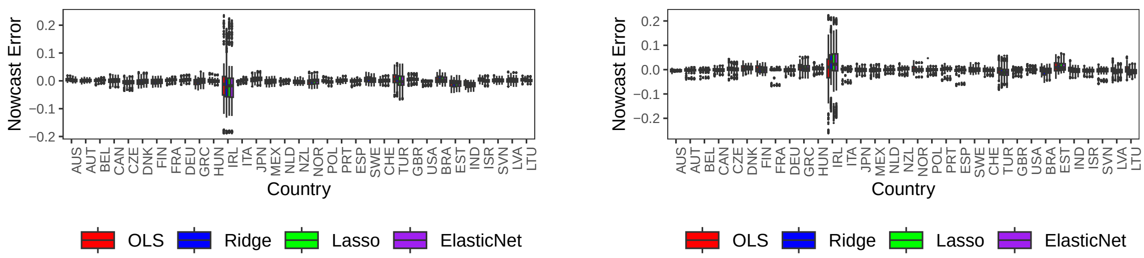

Note: Average nowcast error over 155 nowcast days and 23 evaluation periods for countries individually. Out-of-sample split = 30%.

Figure 6: Distribution of nowcast errors across countries for target quarter-on-quarter GDP growth

(a) Pooled models

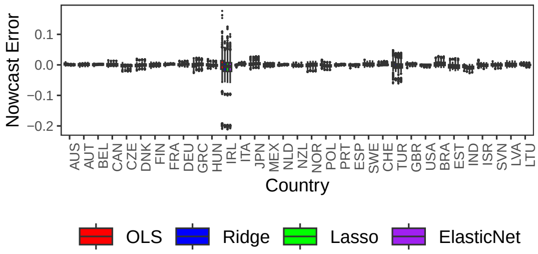

(b) Non-pooled models

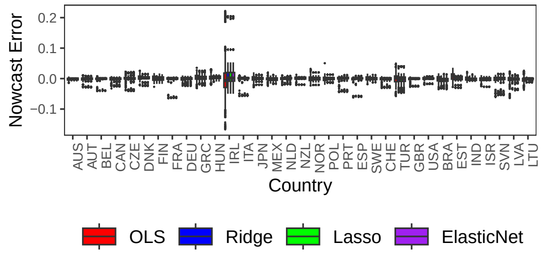

Note: Please see the footnote of Figure 5.

4.4 Shrinkage parameter lambda

Figure 7 and 8 show the average cross-country amount of coefficient shrinkage for the full nowcast period of 155 days for all penalised regression models. This is expressed by value lambda, which is automatically selected by means of cross-validation. Higher values of lambda indicate stronger regularisation and thus, shrinkage of the coefficients, which means that they will become lower. It can be observed that Ridge has higher average lambda values than Lasso and Elastic Net over time. This can be explained by its theoretical background. Lasso sets many coefficients to 0, and Ridge still includes all predictors, which means that even with large lambda values, Ridge can still retain important information from all predictors. If Lasso would have high values of lambda, coefficients of important predictors that are selected may turn out too low, where reducing model complexity potentially leads to underfitting. For pooled models Lasso and Elastic Net in figure 7b and 8b, mean lambda values slowly increase when approaching the end of the nowcasting period, for both yearly and quarterly data, possibly due to staggered release of information during the nowcast period. The lambda values increase to prevent overfitting to the newly available data. The mean lambda for Ridge shows high volatility in all figures, which may be explained by the constant adjustments for all predictors to shrink them proportionally. Further research on this will be conducted with Figure 9.

Figure 7: Mean lambdas over time for target year-on-year GDP growth

(a) Non-pooled models

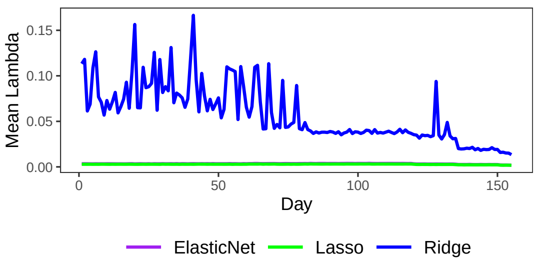

(b) Pooled models

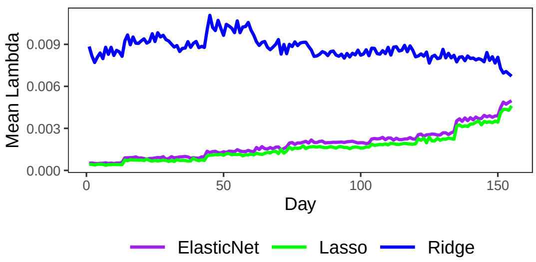

Note: Cross-country shrinkage parameter lambda over the full nowcast period of 155 days for all penalised regression models. Out-of-sample split = 30%.

Figure 8: Mean lambdas over time for target quarter-on-quarter GDP growth

(a) Non-pooled models

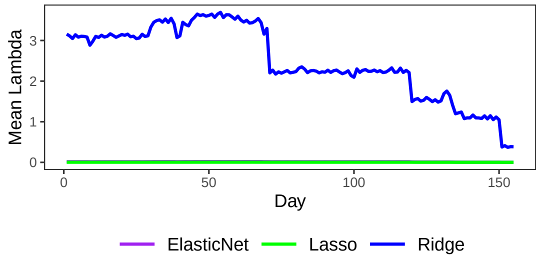

(b) Pooled models

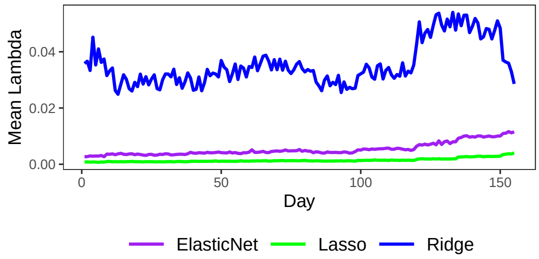

Note: Please see the footnote of Figure 7.

Further analysis of the distribution of shrinkage parameter lambda for Ridge regression is necessary to find the underlying reasons for the higher average values and volatility. Figure 9 demonstrates the distribution of lambda values of the non-pooled models for individual countries. For non-pooled models, the Ridge regression models are estimated separately for all countries by means of cross-validation for all nowcast days of every evaluation period. Figure 9a demonstrates that most countries have low shrinkage parameters for year-on-year GDP growth, with values close to to 0. However, Israel (IRL), Poland (POL) and Indonesia (IND) have a high number outliers for lambda, with values approaching 50. Outliers for Israel can be explained by more its volatile GDP growth rates that differ slightly from the other countries, previously displayed in figure 1. For these countries with high values of lambda, more shrinkage of the coefficients may be necessary for accurate predictions and lower nowcast errors. Figure 9b shows closer distributions of lambda values across different countries, with less outliers. Some countries, such as Estonia (EST), Lithuania (LTU), and Israel (IRL), exhibit distributions of higher lambda values, indicating a stronger regularisation effect for the coefficients of these countries.

Figure 9: Ridge regression distribution of lambdas for individual countries

(a) Target year-on-year GDP growth

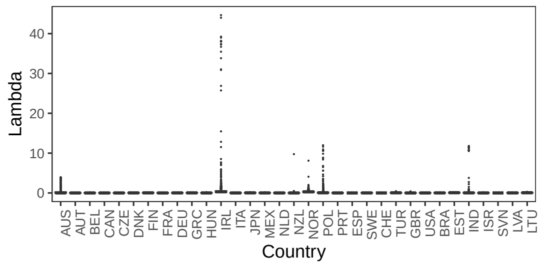

(b) Target quarter-on-quarter GDP growth

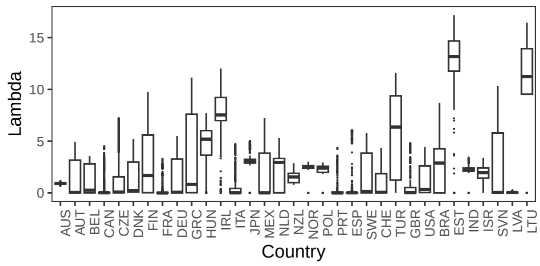

Note: Distributions of shrinkage parameter lambda across countries. Out-of-sample split = 30%.

4.5 Predictor importance

Figure 10 shows the count of selected variables in predictions by Lasso for year-on-year and quarter-on-quarter GDP. Figures 10a and 10b confirm existing theory about Lasso that there is a relation between the number of times predictors that are included in the model and the size of shrinkage parameter lambda (van Wieringen, 2018). The amount of penalisation determines the selection frequency, where less predictors are allowed for higher values of log lambda. This a result of stronger regularisation that lowers the coefficients and will decrease many to 0.0, thus excluding them from the model. There are 58 predictors in the pooled Lasso model, and the selection count has a maximum value of 40 for Y-O-Y, and 30 for Q-O-Q GDP growth, which means that Lasso discards at least 15 predictors for both data sets. As Ridge outperforms Lasso in quarter-on-quarter GDP growth, we can infer that Lasso might be overly enthusiastic in eliminating predictors by setting too many coefficients to zero, including those that are actually relevant. This strict elimination can lead to underfitting, reducing Lasso's predictive accuracy.

Figure 10: Pooled Lasso selection count of predictors for log lambda values

(a) Target year-on-year **(b)** Target quarter-on-quarter

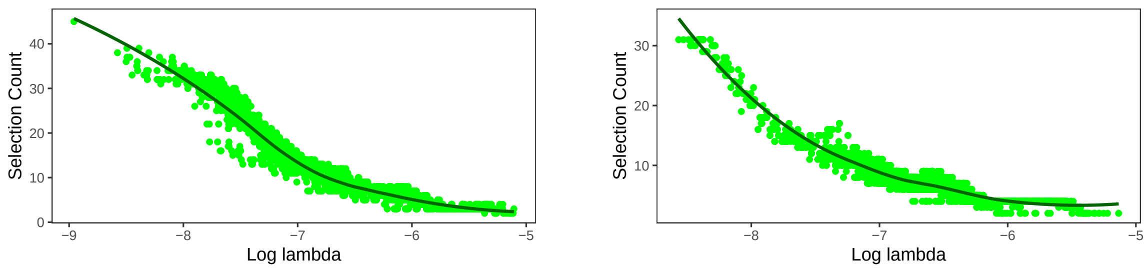

Note: Average count of selected predictors where the coefficients are not equal to zero, plotted against the shrinkage parameter log lambda values. Out-of-sample split = 30%.

Figures 11 and 12 demonstrate for every predictor how often they are selected for all pooled and non-pooled models. We can see that OLS and Ridge consistently select all predictors, indicating that there is no variable selection. Ridge is not able to shrink coefficients to zero, and OLS finds a linear relationship between GDP growth and all predictors. Lasso definitely shows that it is able to eliminate variables, however, in figure 11a and Figure 11b are selected at least once in all regressions. The predictors that are most often selected for non-pooled models are one lag of GDP, and the third lag of BSM and the average BSM. Elastic Net shows moderate predictor elimination, with percentages being slightly higher than Lasso always. We can conclude from this that no matter the relative weight given to L1 or L2, Elastic Net behaves more like Lasso than Ridge regarding variable selection. This is because the weights (α) for Y-O-Y and Q-O-Q GDP growth differ, namely 0.8 and 0.2, respectively. The intercept is always included for all models, shown by the selection count of 100%.

Figure 11: Selection count of predictors for all non-pooled models in

(a) Target year-on-year **(b)** Target quarter-on-quarter

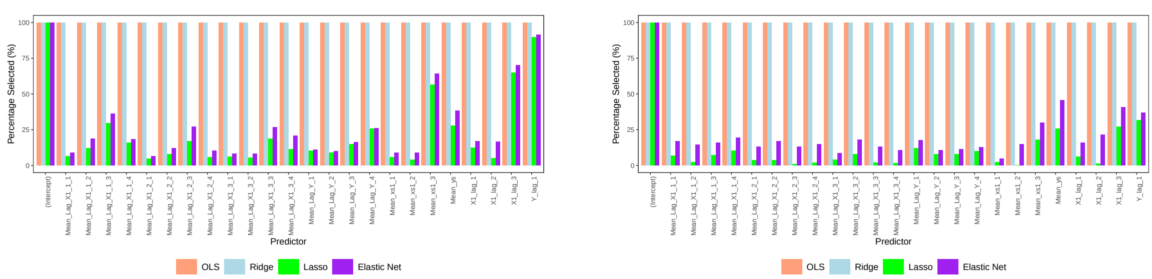

Note: Percentage count of how often predictors are selected for all 25 predictors of non pooled models. Out-of-sample split = 30%.

Figure 12 shows how often predictors are selected for pooled models. There are significantly more predictors in this figure, which can be explained by the individual-specific effects that are estimated individually for all countries. Intercepts for country 24, 29 and 30 are often selected, corresponding to Türkiye, Indonesia and Israel, respectively. As discussed before, these countries deviate in GDP growth over time compared to others, making their specific intercepts more important for accurate predictions, and thus more often selected. For pooled models there are more important predictors in addition to the ones mentioned before, such as the first lag of the average BSM and GDP.

Figure 12: Selection count of predictors for all pooled models

(a) Target year-on-year **(b)** Target quarter-on-quarter

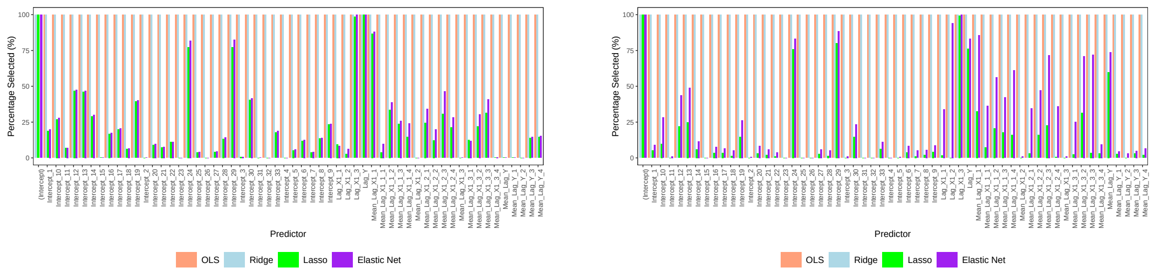

Note: Percentage count of how often predictors are selected for all 58 predictors of pooled models. Out-of-sample split = 30%.

Figures 13 and 14 depict the relative mean coefficient size of all predictors for the models to further dive into predictor importance. For every model, the relative coefficient size is calculated by dividing the coefficients by the highest coefficient given to a predictor for that model. OLS indicates that no shrinkage is applied, because it consistently shows the highest relative coefficients for most predictors. Ridge shows smaller coefficients compared to OLS, but still maintains high relative sizes for most predictors. Lasso shrinks many coefficients to zero, and Elastic Net has lower coefficients than Ridge, but higher than Lasso, indicating its balance in between methods. The most important predictors are lags of GDP and mean values of those lags.

Figure 13: Relative coefficient size of predictors for all non-pooled models

(a) Target year-on-year **(b)** Target quarter-on-quarter

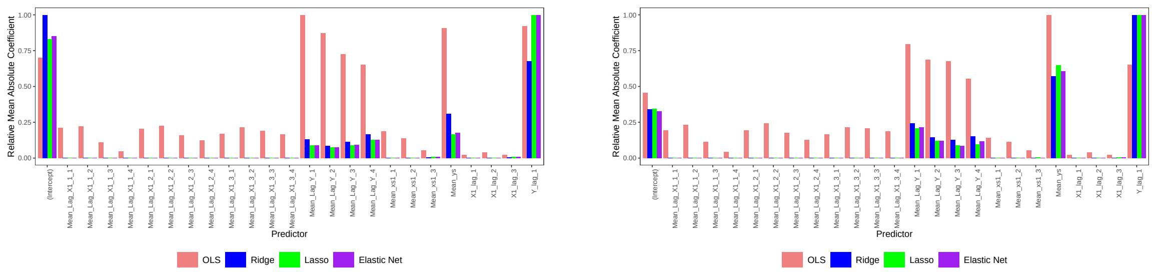

Figure 14: Relative coefficient size of predictors for all pooled models

(a) Target year-on-year

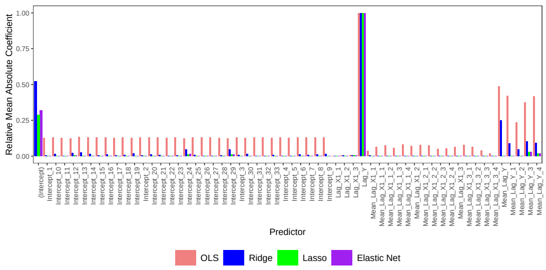

(b) Target quarter-on-quarter

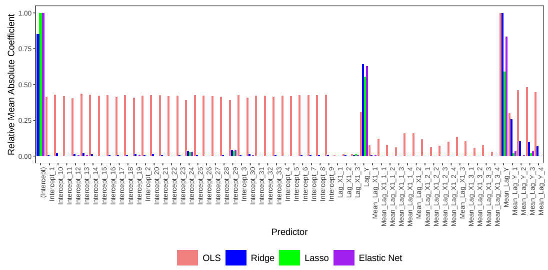

4.6 Volatility and nowcast errors

We want to generalise the predictive accuracy of the models to other datasets, by plotting the mean absolute nowcast errors to the three period rolling window standard deviation of the target variable GDP in Figure 15. Figure 15a shows that Lasso and Elastic Net perform best for Y-O-Y, where the lines are below the others in general, indicating that its mean absolute nowcast errors are lowest for different values of standard deviations. Figure 15b shows that Ridge and Elastic Net outperform the other models. We recommend to use Elastic Net for new datasets as this estimation method performs best in general. OLS is definitely outperformed by Elastic Net, and the differences between Elastic Net and Lasso and Ridge are negligible.

Figure 15: Rolling standard deviation of actual values to mean absolute nowcast errors for pooled models

(a) Target year-on-year

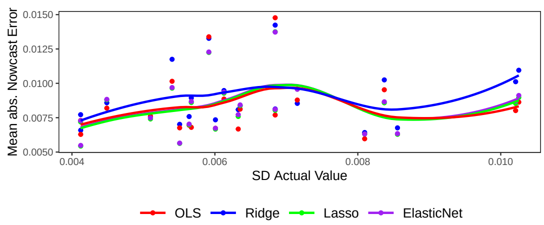

(b) Target quarter-on-quarter

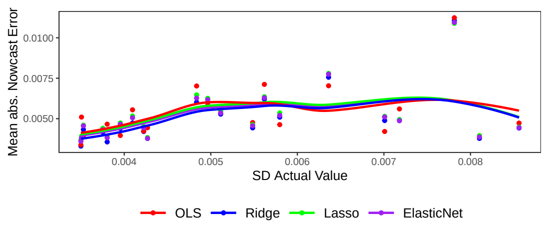

Note: Rolling standard deviation of GDP growth over 3 periods to indicate the relationship between volatility of the target variable and nowcast errors. Out-of-sample split = 30%.

4.7 Robustness to additional predictors

Figure 16 presents the nowcast average root mean squared forecast error for additional variables for the best performing models, pooled Lasso for yearly GDP growth, and pooled Ridge for quarterly GDP growth. The predictor variables used are Business Survey Manufacturing, Business Survey Services and Industrial Production and combinations of them. Figure 4b shows that for Y-O-Y GDP growth the best performing variable is industrial production, as it has the lowest RMSFE over the full nowcast period. Business survey Services also predicts accurately, and the combination of BSS and IP together is almost as good. Adding a third predictor variable decreases the accuracy of predictions. Figure 16a demonstrates that for Q-O-Q GDP other predictor variables are better, namely BSM and BSS. However, the other combinations of predictors perform quite similar and the differences are very small. These differences can possibly be be attributed to their economic relevance that may better align with yearly or quarterly GDP growth. IP may be less volatile and better capture long term trends which makes it more suitable for yearly GDP. It is possible that BSM and BSS better capture short-term variations, which makes them better predictors for quarterly GDP.

Figure 16: Nowcast average RMSFE three predictors for targets year-on-year and quarter-on-quarter

(a) Y-O-Y RMSFE for pooled Lasso **(b)** Q-O-Q RMSFE for pooled Ridge

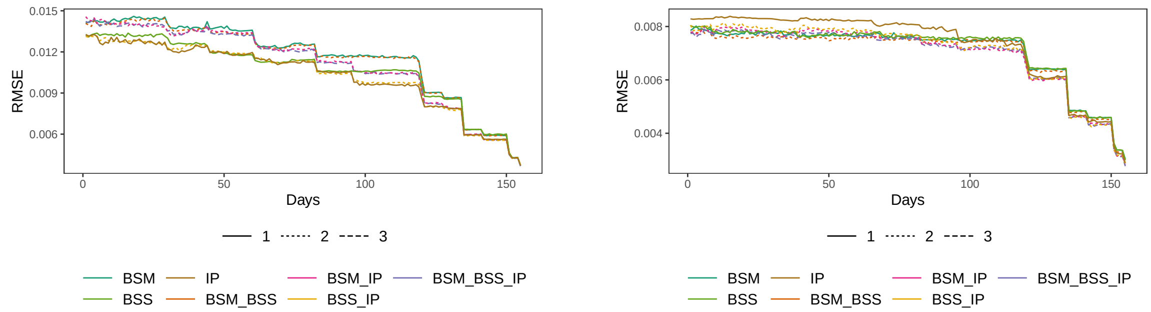

Note: Root mean squared forecast error (RMSFE) for full nowcast period with different combinations of predictors BSM, BSS and IP. Out-of-sample split = 30%.

5 Discussion and Conclusion

In this thesis, the goal was to improve the forecasting of GDP growth by incorporating regularisation methods for shrinkage and lag selection. Building on the panel mixed data sampling (PMIDAS) approach introduced by Fosten and Nandi (2023), this study applied penalised regression methods Ridge, LASSO, and Elastic Net to analyse real GDP growth with cross-country dependencies. Data from the OECD's Main Economic Indicators (MEI) database is analysed, covering a large set countries and featuring various key predictors, spanning a period from 2001 to 2020. The predictor variables included business surveys on Manufacturing and Services, as well as the Industrial Production index. The objective was to achieve more accurate predictions compared to traditional Ordinary Least Squares and time series AR methods, measured by the root mean squared forecast error (RMSFE). The results presented in this thesis show that estimating PMIDAS with penalised regression methods instead of OLS indeed yields more accurate predictions of GDP growth, demonstrated by lower RMSFE values over the full nowcast period of 155 days. The regularisation methods introduce a bias-variance trade-off, which reduces overfitting and improves generalisation by adding a penalty to the coefficient magnitudes.

The first main conclusion of this thesis is that pooled models consistently have lower RMSFE when compared to non-pooled models, which can be explained by the lower number of estimated parameters, a larger sample size, and reduced variance. Data frequency has a high impact on model performance. Models with the highest prediction accuracy differ for quarterly and yearly GDP growth. The average level of RMSFE for Q-O-Q GDP growth of all models is significantly lower than for Y-O-Y GDP growth, possibly because Q-O-Q GDP growth is less volatile and easier to predict. For yearly GDP growth, pooled Lasso and Elastic Net ($\alpha = 0.8$) demonstrate superior predictive accuracy over non-pooled models and AR, particularly when using Industrial

Production (IP) as the main predictor. Conversely, for quarterly GDP growth, pooled Ridge and Elastic Net ($\alpha = 0.2$) outperform the other estimation methods, with Business Survey Manufacturing (BSM) and Business Survey Services (BSS) proving to be the most effective predictors. These results are robust to other out-of-sample split ratios, as they are consistent across two other sample splits of training and test data. In addition, higher out-of-sample split ratios result in a higher level of RMSFE. We find that countries Israel and Türkiye are most difficult to predict, as they exhibit the largest distribution of nowcast errors. In addition, we find that Ridge shrinks coefficients the most, as this model has the highest average lambda values over time. Pooled Lasso and Elastic Net shrink coefficients much less, with a slight increase towards the end of the nowcast period, possibly due to the staggered release of information. We find a significant relation between the shrinkage parameter and the number of selected variables for Pooled Lasso, where less predictors are allowed for higher values of log lambda. For Q-O-Q GDP growth strong shrinkage of Lasso led to underfitting of the model and lower predictive accuracy. Furthermore, based on the selection count and relative coefficient size, lags one and three of GDP and the predictor variable BSM, as well as their averages, appear to be the most important in predicting GDP growth for both frequencies. Finally, based on the effect of the rolling standard deviation of target variable GDP on the nowcast error, we recommend to use Elastic Net for new data sets. The differences between Elastic Net and the other penalised regression methods in prediction accuracy are negligible, and we find that Elastic Net is more flexible as it combines both penalties of the coefficients. Elastic Net combines the strengths of both models, but more importantly, it addresses their limitations.

Despite the overall success of applying penalised regression methods for nowcasting, this thesis has certain limitations, such as high computational complexity and very long run times for all machine learning models. In addition, these models lack interpretability regarding variable and lag selection, making it difficult to understand why they have better prediction performance. There is a risk of overfitting the training data and some model assumptions such as possible cross-sectional dependence and homogeneous slope coefficients across all countries for pooled models may not hold true. Finally, to be able to make a full comparison with the PMIDAS model of Fosten and Nandi (2023), we were bound to their selection of predictor variables, but it is possible that these predictors did not capture all relevant economic dynamics to make accurate forecasts of real GDP growth.

For future research, we recommend conducting a sensitivity analysis to better understand how sensitive the nowcast performance is to different model specifications such as lag choice, frequency of variables, and sample periods. Furthermore, we suggest looking into other machine learning models such as Neural Networks and Random Forests. In addition, it may be relevant to add global economic indicators to capture cross-country interdependencies and investigate spillover effects between the countries.

6 Acknowledgements

I would like to thank Dr. A. Pick from Erasmus School of Economics, Erasmus University Rotterdam for his support and guidance during the process of writing this thesis.

References

Babii, A., Ball, R. T., Ghysels, E. & Striaukas, J. (2020, January). Machine Learning Panel Data Regressions with an Application to Nowcasting Price Earnings Ratios. *Social Science Research Network*. doi: 10.2139/ssrn.3670847

Baltagi, B. H. (2008, January). Forecasting with panel data. *Journal of forecasting*, *27*(2), 153–173. doi: 10.1002/for.1047

Beyhum, J. & Striaukas, J. (2023, December). Sparse plus dense MIDAS regressions and nowcasting during the COVID pandemic.

Blei, D. M. (2015, December). *Regularized regression* (Tech. Rep.). Columbia University. Retrieved from `https://www.cs.columbia.edu/~blei/fogm/2015F/notes/regularized-regression.pdf`

Business tendency and consumer opinion surveys. (2024). Retrieved from `https://doi.org/10.1787/data-00041-en` (In Main Economic Indicators)

Chudik, A. & Pesaran, M. H. (2015, October). Common correlated effects estimation of heterogeneous dynamic panel data models with weakly exogenous regressors. *Journal of econometrics*, *188*(2), 393–420. doi: 10.1016/j.jeconom.2015.03.007

ECB. (2000, October). Different ways of calculating the growth rate of real gdp. *ECB Monthly Bulletin*. Retrieved from `https://www.ecb.europa.eu/pub/pdf/other/mb200010_focus02.en.pdf`

Foroni, C., Marcellino, M. & Schumacher, C. (2015, January). Unrestricted Mixed Data Sampling (MIDAS): MIDAS Regressions with Unrestricted Lag Polynomials. *Journal of the Royal Statistical Society Series A: Statistics in society*, *178*(1), 57–82. doi: 10.1111/rssa.12043

Fosten, J. & Nandi, S. (2023). Nowcasting from cross-sectionally dependent panels. *Journal of Applied Econometrics*, *38*, 898–919. doi: 10.1002/jae.2980

Ghysels, E. (2016). Macroeconomics and the reality of mixed frequency data. *Journal of Econometrics*, *193*(2), 294-314. doi: 10.1016/j.jeconom.2016.04.008

Hastie, T., Qian, J. & Tay, K. (2023, March). *An Introduction to 'glmnet'.* Retrieved from `https://glmnet.stanford.edu/articles/glmnet.html`

Hastie, T., Tibshirani, R. J. & Friedman, J. (2013). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Retrieved from `http://catalog.lib.kyushu-u.ac.jp/ja/recordID/1416361`

International Monetary Fund. (2015, September). *Israel: Concluding Statement of the 2015 Article IV Consultation.* Retrieved from `https://www.imf.org/en/News/Articles/2015/09/28/04/52/mcs062415`

Kant, D., Pick, A. & Winter, J. d. (2022, January). Nowcasting GDP using machine learning methods. *Social Science Research Network*. doi: 10.2139/ssrn.4276247

Liu, H., Liu, Y., Li, G. & Wen, L. (2021, July). Tourism demand nowcasting using a LASSO-MIDAS model. *International journal of contemporary hospitality management*, *33*(6), 1922–1949. doi: 10.1108/ijchm-06-2020-0589

Richardson, A., Mulder, T. & Vehbi, T. L. (2018, January). Nowcasting New Zealand GDP using machine learning algorithms. *Social Science Research Network*. doi: 10.2139/ssrn.3256578

Siliverstovs, B. (2016, August). Short-term forecasting with mixed-frequency data: a MIDASSO approach. *Applied economics*, *49*(13), 1326–1343. doi: 10.1080/00036846.2016.1217310

Tiffin, A. (2016, January). Seeing in the Dark: A Machine-Learning Approach to Nowcasting in Lebanon. *Social Science Research Network*. doi: 10.2139/ssrn.2770291

van Wieringen, W. (2018). *Lasso regression* (Tech. Rep.). Department of Epidemiology and Biostatistics, VUmc and Department of Mathematics, VU University. Retrieved from `\url{https://www.few.vu.nl/~wvanwie/Courses/HighdimensionalDataAnalysis/WNvanWieringen_HDDA_Lecture56_LassoRegression_20182019.pdf}`

Wang, W., Zhang, X. & Paap, R. (2019, March). To pool or not to pool: What is a good strategy for parameter estimation and forecasting in panel regressions? *Journal of applied econometrics*, *34*(5), 724–745. doi: 10.1002/jae.2696

Zhang, Q., Ni, H. & Hao, X. (2023, May). Nowcasting Chinese GDP in a data-rich environment: Lessons from machine learning algorithms. *Economic modelling*, *122*. doi: 10.1016/j.econmod.2023.106204

7 Appendix

A Data

Table 5: Predictor variables availability and country coverage

Abbr.	**Variable**	**Available time period**	**N**	**Country names**
BSM	Manufacturing (Business surveys)	2001 - 2020	34	Australia, Austria, Belgium, Canada, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Japan, Luxembourg, Mexico, Netherlands, New Zealand, Norway, Poland, Portugal, Spain, Sweden, Switzerland, Turkey, United Kingdom, United States, Brazil, Estonia, India, Israel, Slovenia, Latvia, Lithuania
BSS	Services (Business surveys)	2003 - 2020	21	Austria, Belgium, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Netherlands, Poland, Portugal, Spain, Sweden, United Kingdom, Estonia, Slovenia, Latvia, Lithuania
IP	Industrial production	2001 - 2020	35	Austria, Belgium, Canada, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Japan, Korea, Luxembourg, Mexico, Netherlands, Norway, Poland, Portugal, Slovak Republic, Spain, Sweden, Turkey, United Kingdom, United States, Brazil, Chile, Estonia, India, Israel, Slovenia, Latvia, Lithuania, Costa Rica
CON	Consumer	2001 - 2023	19	Austria, Belgium, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Italy, Latvia, Lithuania, Netherlands, Poland, Portugal, Slovenia, Spain, Sweden
TRA	Trade	1999 - 2021	19	Austria, Belgium, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Italy, Latvia, Lithuania, Netherlands, Poland, Portugal, Slovenia, Spain, Sweden
CST	Construction	2002 - 2020	19	Austria, Belgium, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Italy, Latvia, Lithuania, Netherlands, Poland, Portugal, Slovenia, Spain, Sweden

Note: This table lists the available time periods and countries in the sample for all predictor variables.

Table 6: Cross-Country Correlation Matrix for Y-O-Y data

	GDP	IP	Consumer	Services	Trade	Manufacturing	Construction
GDP	1.000						
IP	-0.054	1.000					
Consumer	0.064	0.143	1.000				
Services	-0.174	0.285	-0.084	1.000			
Trade	0.013	0.177	0.404	-0.005	1.000		
Manufacturing	-0.164	0.334	-0.040	0.955	0.037	1.000	
Construction	-0.069	0.408	0.306	0.249	0.452	0.322	1.000

Figure 17: Distribution cross-country of GDP and predictors over time

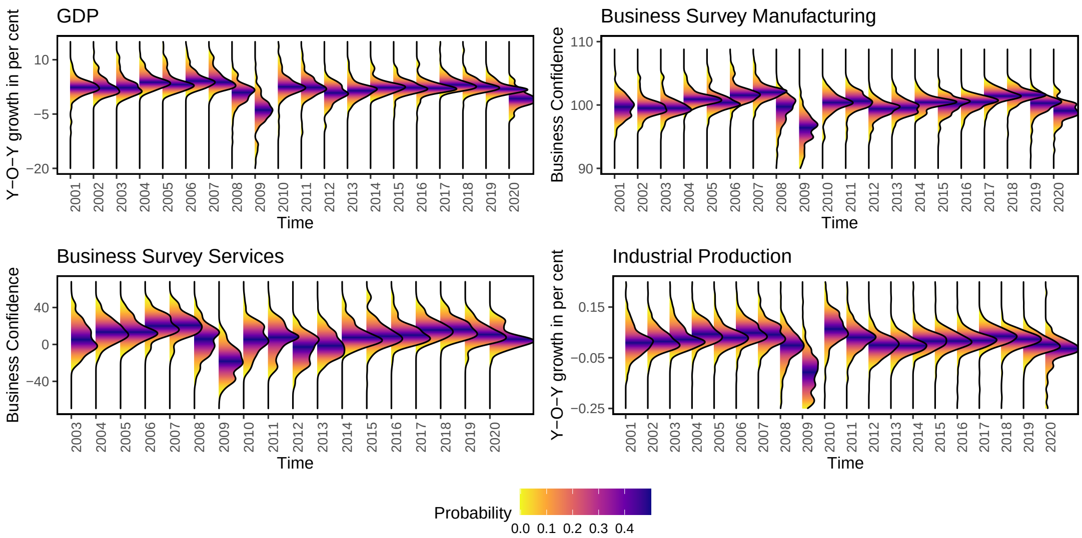

Note: This figure shows the cross-country probability distribution of GDP, BSM, BSS and IP from 2001 to 2020. Darker colours indicate higher probability, as shown in the legend.

Table 7: VIF values and cross-correlation matrix for transformed predictors of GDP and BSM

(a) Part 1

	L_Y	L_X1_1	L_X1_2	L_X1_3	M_L_Y	M_L_X1_1	M_L_X1_2	M_L_X1_3	M_L_Y_1	M_L_Y_2	M_L_Y_3	M_L_Y_4	M_L_X1_1_1
VIF	2.121	104.447	361.471	115.561	45.442	1437.214	11469.598	20806.794	83.656	83.475	71.566	29.188	23493.368
L_Y	1.000	0.677	0.656	0.631	0.602	0.519	0.542	0.556	0.549	0.424	0.273	0.128	0.560
L_X1_1	0.677	1.000	0.984	0.944	0.657	0.678	0.702	0.710	0.579	0.442	0.286	0.130	0.703
L_X1_2	0.656	0.984	1.000	0.985	0.641	0.701	0.709	0.701	0.537	0.390	0.232	0.083	0.679
L_X1_3	0.631	0.944	0.985	1.000	0.613	0.709	0.701	0.678	0.490	0.338	0.180	0.040	0.644
M_L_Y	0.602	0.657	0.641	0.613	1.000	0.862	0.901	0.924	0.912	0.706	0.454	0.214	0.931
M_L_X1_1	0.519	0.678	0.701	0.709	0.862	1.000	0.989	0.956	0.687	0.472	0.250	0.052	0.906
M_L_X1_2	0.542	0.702	0.709	0.701	0.901	0.989	1.000	0.989	0.753	0.546	0.323	0.112	0.957
M_L_X1_3	0.556	0.710	0.701	0.678	0.924	0.956	0.989	1.000	0.811	0.617	0.398	0.179	0.989
M_L_Y_1	0.549	0.579	0.537	0.490	0.912	0.687	0.753	0.811	1.000	0.912	0.705	0.453	0.862
M_L_Y_2	0.424	0.442	0.390	0.338	0.706	0.472	0.546	0.617	0.912	1.000	0.912	0.704	0.688
M_L_Y_3	0.273	0.286	0.232	0.180	0.454	0.250	0.323	0.398	0.705	0.912	1.000	0.911	0.472
M_L_Y_4	0.128	0.130	0.083	0.040	0.214	0.052	0.112	0.179	0.453	0.704	0.911	1.000	0.248

(b) Part 2

	M_L_X1_1_2	M_L_X1_1_3	M_L_X1_1_4	M_L_X1_2_1	M_L_X1_2_2	M_L_X1_2_3	M_L_X1_2_4	M_L_X1_3_1	M_L_X1_3_2	M_L_X1_3_3	M_L_X1_3_4
VIF	27183.842	30310.615	24984.649	27363.490	33267.669	32119.878	12739.417	27104.931	32905.187	31339.854	1747.413
M_L_X1_1_1	0.907	0.691	0.433	0.989	0.844	0.608	0.344	0.956	0.770	0.521	0.256
M_L_X1_1_2	1.000	0.907	0.692	0.957	0.989	0.844	0.607	0.989	0.956	0.771	0.518
M_L_X1_1_3	0.907	1.000	0.905	0.772	0.957	0.989	0.840	0.843	0.989	0.956	0.765
M_L_X1_1_4	0.692	0.905	1.000	0.521	0.772	0.956	0.988	0.607	0.842	0.989	0.954
M_L_X1_2_1	0.989	0.957	0.772	1.000	0.908	0.693	0.432	0.989	0.844	0.608	0.342
M_L_X1_2_2	0.844	0.989	0.957	0.908	1.000	0.908	0.691	0.957	0.989	0.845	0.604
M_L_X1_2_3	0.608	0.844	0.989	0.693	0.908	1.000	0.905	0.771	0.957	0.989	0.839
M_L_X1_2_4	0.344	0.607	0.840	0.432	0.691	0.905	1.000	0.520	0.768	0.955	0.988
M_L_X1_3_1	0.956	0.989	0.843	0.989	0.957	0.771	0.520	1.000	0.906	0.691	0.430
M_L_X1_3_2	0.770	0.956	0.989	0.844	0.989	0.957	0.768	0.906	1.000	0.907	0.685
M_L_X1_3_3	0.521	0.771	0.956	0.989	0.608	0.845	0.989	0.691	0.907	1.000	0.903
M_L_X1_3_4	0.256	0.518	0.765	0.954	0.342	0.604	0.839	0.430	0.685	0.903	1.000

Note: Cross-correlations for Y-O-Y real GDP growth and Business Survey Manufacturing.

B Results

Figure 18: Elastic Net nowcast RMSFE for alpha parameters for target year-on-year

(a) Out-of-sample split 20%

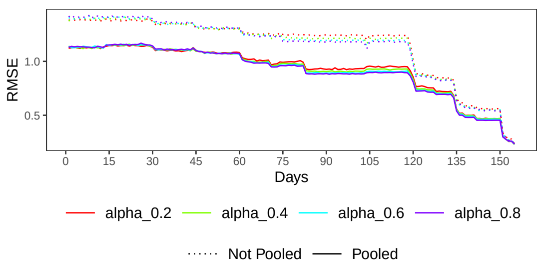

(b) Out-of-sample split 40%

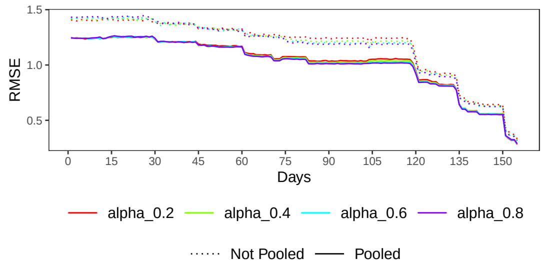

Note: This figure plots for each alpha ($\alpha = 0.2, 0.4, 0.6, 0.8$) of Elastic Net the cross-country average root mean squared forecast error (RMSFE) on the full nowcast period of 155 days. The predictor variable is Business Survey Manufacturing, and the target variable is year-on-year GDP, with out-of-sample splits of 20% and 40%.

Figure 19: Elastic Net nowcast RMSFE for alpha parameters for target quarter-on-quarter

(a) Out-of-sample split 20%

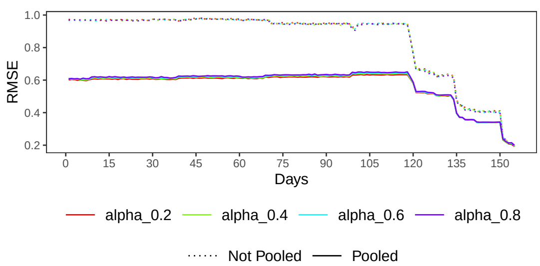

(b) Out-of-sample split 40%

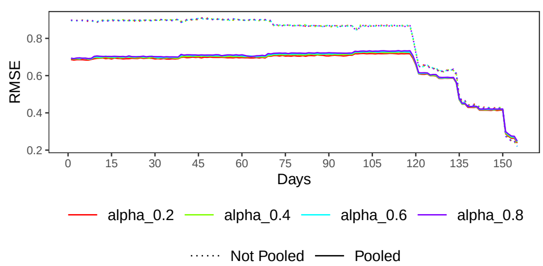

Note: This figure plots for each alpha ($\alpha = 0.2, 0.4, 0.6, 0.8$) of Elastic Net the cross-country average root mean squared forecast error (RMSFE) on the full nowcast period of 155 days. The predictor variable is Business Survey Manufacturing, and the target variable is quarter-on-quarter GDP, with out-of-sample splits of 20% and 40%.

Figure 20: Nowcast average RMSFE of estimation methods for target year-on-year

(a) Out-of-sample split 20%

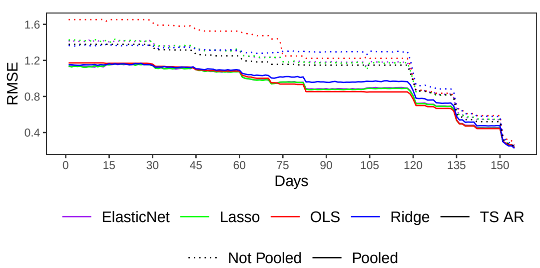

(b) Out-of-sample split 40%

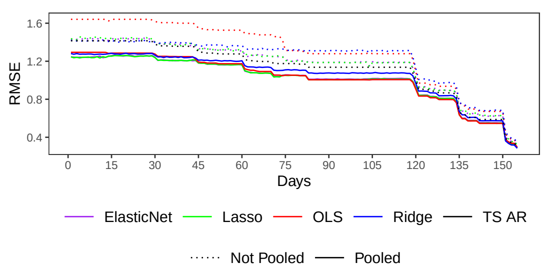

Note: Root mean squared forecast error (RMSFE) for full nowcast period.

Figure 21: Nowcast average RMSFE of estimation methods for target quarter-on-quarter

(a) Out-of-sample split 20%   **(b)** Out-of-sample split 40%

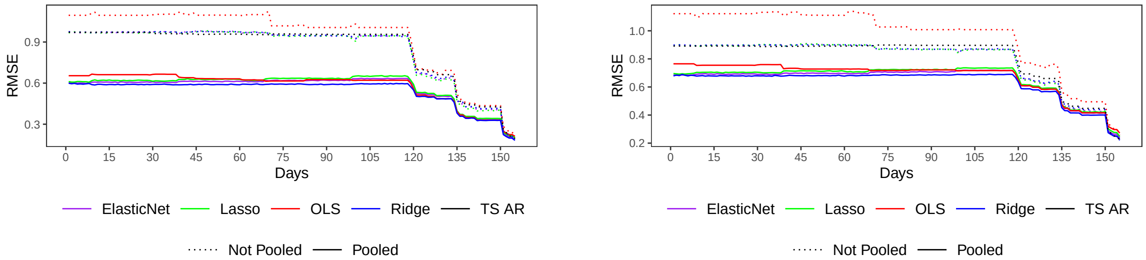

Note: Root mean squared forecast error (RMSFE) for full nowcast period.

Table 9: Average RMSFE for target year-on-year of estimation methods for various out-of-sample split ratios

Estimation method		Split		
Method	**Pooled**	**0.2**	**0.3**	**0.4**
Lasso	Pooled	0.0094	0.0104	0.0105
OLS	Pooled	0.0094	0.0105	0.0106
Elastic Net	Pooled	0.0095	0.0105	0.0105
Ridge	Pooled	0.0101	0.0113	0.0112
AR	AR	0.0110	0.0116	0.0113
Lasso	Not Pooled	0.0114	0.0121	0.0117
Elastic Net	Not Pooled	0.0114	0.0121	0.0117
Ridge	Not Pooled	0.0123	0.0132	0.0128
OLS	Not Pooled	0.0126	0.0131	0.0130

Table 10: Average RMSFE for target quarter-on-quarter of estimation methods for various out-of-sample split ratios

Estimation method		Split		
Method	**Pooled**	**0.2**	**0.3**	**0.4**
Ridge	Pooled	0.0059	0.0068	0.0068
Elastic Net	Pooled	0.0061	0.0069	0.0069
OLS	Pooled	0.0062	0.0071	0.0072
Lasso	Pooled	0.0062	0.0070	0.0070
Lasso	Not Pooled	0.0085	0.0084	0.0080
AR	AR	0.0086	0.0085	0.0081
Elastic Net	Not Pooled	0.0086	0.0085	0.0080
Ridge	Not Pooled	0.0087	0.0086	0.0082
OLS	Not Pooled	0.0094	0.0096	0.0096