

Long Term Violence Report

Data Analytics Lab

July 2021

Table of Contents

1	Introduction.....	3
2	Main Findings.....	4
3	Potential Approaches.....	6
4	Literature.....	7
5	Data.....	8
5.1	Local Authority Data	8
5.2	Macroeconomic Data.....	22
5.3	Crimes	24
6	Modelling.....	26
6.1	ARIMA, ETS and Auto-Regressive Distributed Lag modelling	26
6.2	Multivariate Models.....	28
6.3	Cross-Validation.....	28
7	Results & Analysis	30
7.1	Univariate Analysis	30
7.1.1	ARIMA and Exponential Smoothing	34
7.2	VAR and ARDL Models.....	36
7.2.1	Vector Auto Regressions	36
7.2.2	Auto-Regressive Distributed Lag Model	46
8	Conclusions & Recommendations.....	58
9	Appendix.....	59
9.1.1	VAR Coefficients	59
9.1.2	Auto Regressive Distributed Lag Coefficients	61
9.1.3	Terms Used	63
10	References.....	64

1 Introduction

The aim of this project is to consider the feasibility of forecasting long term violence (as defined using previous most serious violence projects) using both the standard time series (auto-regressive) properties and using additional information from inside West Midlands Police and open source local and national government data.

The nature of the data involves a number of challenges, some of the data (for example crime and intelligence logs) are daily or more frequent in their nature whereas the local government data is at best monthly and most commonly annual (based on either the fiscal or academic year for a number of series). This necessitates the use of mixed time series methods, merging the data into a monthly data series. This smooths out some of the daily variations and repeats the annual data.

The data were selected to capture the overall feel of the socio-economic situation in the West Midlands area as a whole. Not all the data is based on the whole of the region, rather areas or towns are used as indicative areas. The areas are used to represent the urban areas of the West Midlands and the general environment in the region. The data that is less regular needs to be forecast-able with the predictive interval potentially being problematic if it is too broad.

One of the attractive aspects of pure time series analyses is the fact that the forecasts are based on only that series and no other data is needed and so the forecasts can loop back into the models for later time points. Adding the exogenous factors means that a forecast of these has to be made. This technique is known as *model stacking*; a first model is used for forecasting these factors and this is fed into the second model that forecasts the variable of interest. This approach introduces predictive uncertainty from another source – any univariate predictive model (uses just the time series) would have uncertainty due to variance in the time series of interest, and the model(s) associated with the forecast of the other variables introduces further sources of uncertainty. In light of this, one can use a number of models and compare and/or combine the forecasts to mitigate this uncertainty; but it will remain to some degree.

2 Main Findings

Forecasting a year forward is possible, though the addition of independent variables is not as helpful as one would hope. This is due to the lagging nature of this information and the relatively coarse nature of the data. A better approach is to use a relatively simple Vector Autoregression that accommodates interactions between the regions. A five year estimate is possible, but the prediction intervals are wide, i.e. the amount of uncertainty surrounding the estimates increases the further out in time the forecast is for. In essence this is equivalent to stating that violence in a specific place is going to happen, but the level of this is unsure. The forecasts become rather straight lines, that is picking up a trend but the prediction intervals explode. This type of forecast is of limited use and should not be relied on except in the broadest trends.

An additional factor that has mostly been left out of consideration is the lock down associated with COVID 19. This will have an impact on the figures, however given the sample this was mostly left out of the analyses.

It appears that in this case, keeping the model relatively simple allowing the NPUs to have a knock on effect does give rise to potentially useful insights over 12 months, but to look beyond this is currently too uncertain to be of any significant use.

Below, in Figure 1, is an example of the forecast one year out of the simplest models. There is a degree of variation across the board, but the spreads (measured as prediction intervals) are still substantial. This picks up the observed variation in the data but makes the forecasts less certain. The five year forecasts see these results writ large. The uncertainty associated with stretching the window out such a long way and in light of the extraordinary 18 months leads to considerable prediction intervals as estimated from simulations and purely from the estimated models.

Adding extra information has limited impact as the data too is rather coarse and will introduce more uncertainty into the modeling without adding any significant predictive benefit in a vast majority of areas.

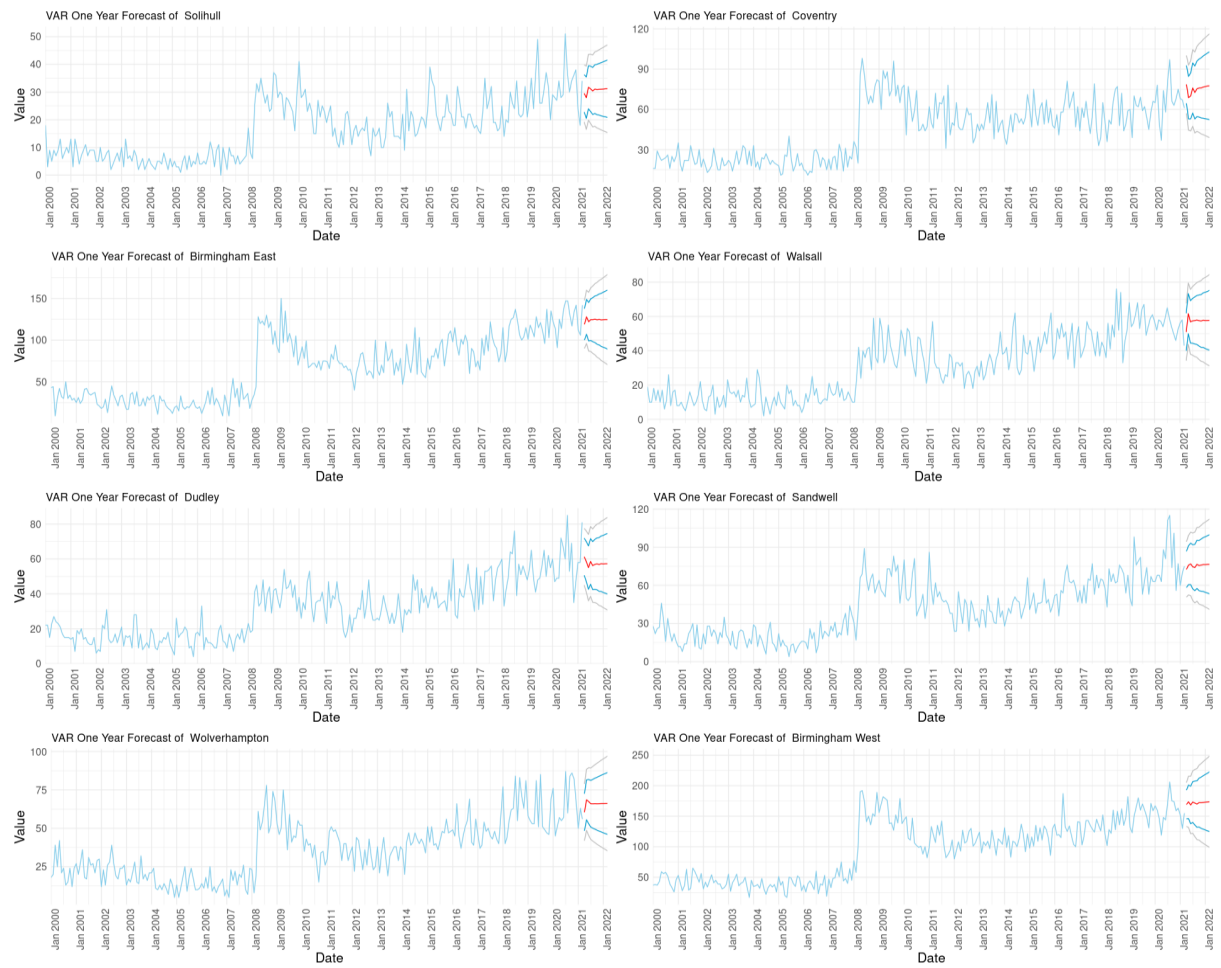


Figure 1 Forecasts of Violent Crime

Key: red mean forecast, sky-blue 80th % interval, grey 95th % interval

3 Potential Approaches

There are a number of methods that might be applicable for the forecasting problem. The diagram below highlights a number of these. The time series aspect is of primary importance and as such the time series models are focused upon. An approach used that amalgamates the various models is highlighted.

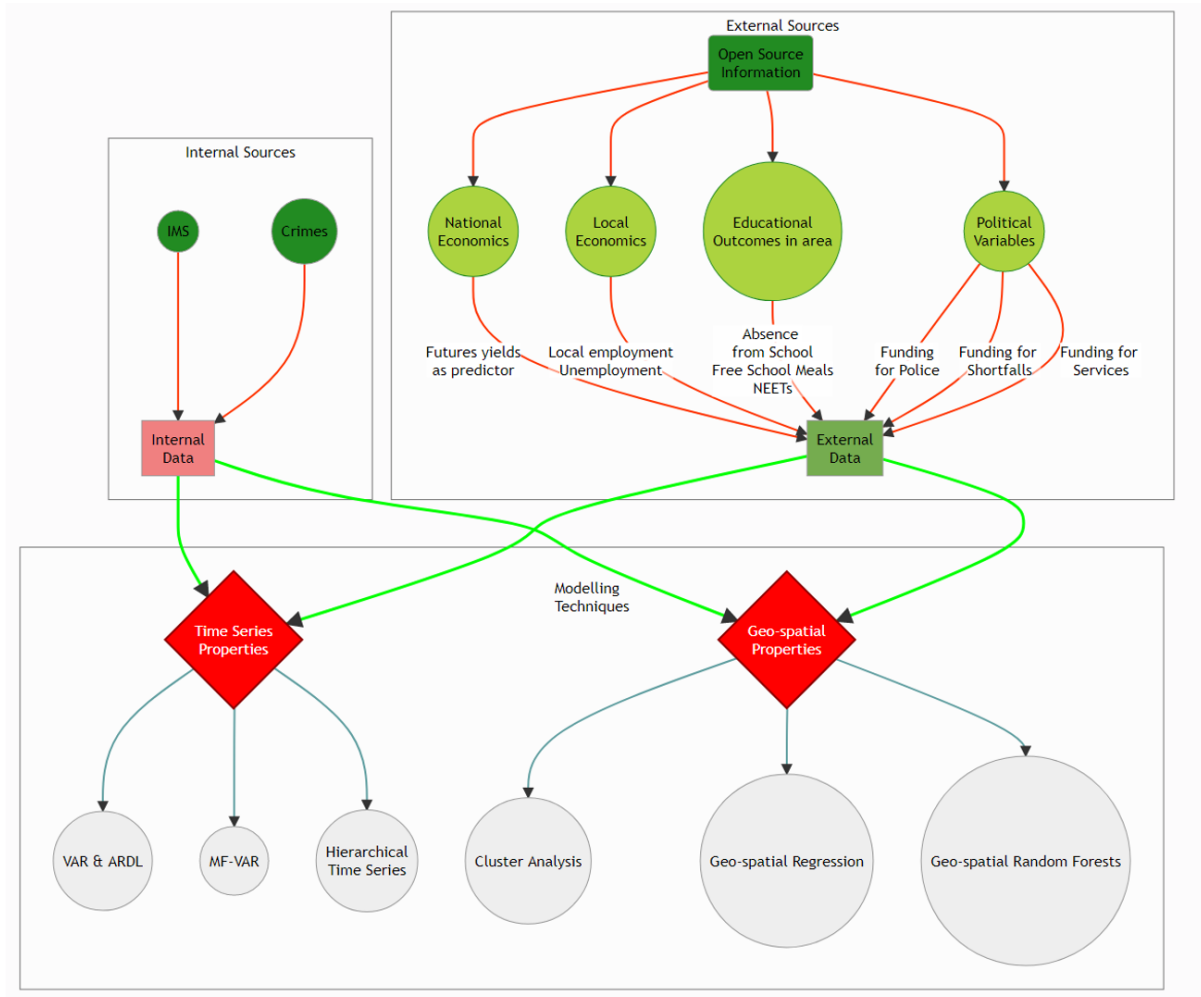


Figure 2 Methods of Data Collection and Modelling Techniques

The Figure above shows the possible approaches available to deal with the types of data in the analysis. A time series approach was taken as this most lends itself to the forecasting of the underlying series, rather than the consideration of other aspects of analysis.

4 Literature

There is some literature where the temporal influences and experiential nature of violence is highlighted (Widom (1989), Widom (1992) with childhood experiences leading increased violence in later life, in addition to other conditioning factors. This is of limited use in this more macro-forecasting approach, though the study does imply that a violent environment does lead to an increase (or at least not a decrease) in violence focusing on the childhood experiences of the perpetrators.

There is some literature concerning the development of the individual into a violent person in the longer-term though this is at the individual level and thus beyond the scope of the work. There is some work (e.g. Stoddard et al. (2015) that looks at factors that are good determinants of the longer term violent behaviours of individuals that can, on a macro-level, be informative of the type of explanatory factors to be considered.

In the Stoddard et al. (2015) study, educational aspirations were seen to be important in the probability of violence in the future. This would suggest that local educational outcomes are important in a particular region. This further raises the potential for macro-economic spending levels to act as a proxy driving educational achievement and aspiration. In order to consider the impact of this, a limited time series and depth (say 6 years) of data might be required. This limited time span would be sufficient to deal with the cohorts of young people going through the system.

Taylor, Ratcliffe, and Perenzin (2015) looks at the prediction of crime in as a long and short term phenomenon with a specific focus on the addition of non-crime data to the crime based data and specifically socio-economic data as proxied by area. Our paper builds on that to include more macro-economic data in the models to capture the zeitgeist of the period in question and the potential direction of the society in those areas and further develops the findings of the meta-analysis of Pratt and Cullen (2005).

This ecological work of Pratt and Cullen (2005) which identifies both racial composition, poverty and unemployment as the main external drivers of crime, with the economic factors generally being more important¹. Interestingly, increased use of the criminal justice system is not found to be as important a factor as these other variables. The belief is that through the choice of variables in the models presented here, elements of the social environment as well as the more macro-economy are included in the forecasts in a novel manner.

Norko and Baranoski (2008) considers the same type of problem from a more sociological & clinical perspective. They discuss the role of substance abuse as well as other factors, pointing out that the economic factors contribute more significantly in the determination of violence than other variables. This US based research all points to a set of additional factors that are important for forecasting crime and violence beyond only the crime statistics.

5 Data

The data used has been extracted from a number of sources. The internal sources used were the Crimes database and tables. These are naturally high frequency, more akin to financial data rather than the annual official statistics found in sources such as NOMIS (2021) and LGInform (2021).

When estimating a predictive model it might in turn be necessary to predict variables other than those in which we are interested. In some cases however, markets exist where it is possible to gain some insights as to the future direction of the economy thus removing the need to forecast such a temperamental animal (or at least where people / organisations think the economy is headed). The market for futures carry the forward looking expectations of market participants. The yield spread, though perhaps not as powerful a tool as it once was is still a useful tool in economic forecasting. Increases in the forward or futures rate spreads are suggestive of economic improvements, increases in GDP and economic activity. This data is available now and is predictive or forward looking by its nature reducing the need to specifically estimate this. Thus the more macro-economic factors such as overall employment can be subsumed into this measure even if it does not deal with local information. The more local data will not be available in such a way, and thus a form of prediction for this localised information is required. This need not be pin-point accurate- e.g. it is not looking at whether Sandwell has a 1% or 1.01% growth¹ in employment, rather the trend is important especially as the window of the forecast increases.

A number of important population centres can act as a bellwether indicator of economic activity in the region as a whole. Thus using economic activity in the local authorities may add value to the model though forecasting might be complex. This data is available from NOMIS via the Labour Market statistics.

5.1 Local Authority Data

Data was acquired from the LGInform (2021) and NOMIS (2021) based on the findings of the likes of Norko and Baranoski (2008) and Pratt and Cullen (2005). This is data collected on an annual basis which requires mixed frequency modelling. There is also an impact of a reduced data set for forecasting; using 10 years of data to predict 5 years hence will be associated with large prediction intervals relative to those predictions with a longer data set.

From this Local Authority data we can see that many areas in the WMP area are facing greater levels of non-working homes than Great Britain as a whole, with only Solihull consistently outperforming Great Britain. The map below (figure 4) presents the median of this statistic and shows the boundaries of the cities as defined by ONS. There is a small area outside the West Midlands, however this is not relevant for the task that these data are used.

¹ It is often said to economists that forecasts were **wrong** as they predicted x% inflation and (x+4 occurred); this is not an exact science and there is sufficient noise in the system to expect substantial deviations, especially with more macro-economic data with its main foibles and problems.

In order to look at other local factors, Nomis and LGINform's data was used to consider demographic and local 'political' dimensions. Data collected based on either the school or financial year (with the relevant dates given as start and end dates) were:

- Proportion of households without work (January starting year)
- Proportion of secondary school students qualifying for free school meals (September starting year)
- School absences in secondary schools (September starting year)
- The proportion of council tax not collected (April starting year)
- The police precept (April starting year)

Each of these seeks to add information about the general areas in the West Midlands. The second and third look at the current education issues in the seven flagged towns. The fourth considers the issues associated with general income levels for the council and the last considers the relative importance of police funding at a local level. The free school meals and absences from school are informed by the Stoddard et al. (2015) study and the economic and social factors are suggested by Norko and Baranoski (2008) and Taylor, Ratcliffe, and Perenzin (2015), where these factors are useful in explaining the crime levels.

The proportion of households without work is obtained via NOMIS and is an indicator of the local economic activity and opportunities available to those in society for whom violent crime might be a choice- if the area has high levels of unemployment the threat of job loss due to the criminal action is not necessarily a dis-incentive to the action. Likewise, the proportion of council tax not collected in a given year is an indicator of the general level of deprivation in the area, again the opportunity costs involved in perpetrating a crime are reduced. The interaction of poverty and education is made through the use of the qualification for free school meals for the older pupils. This is measuring the proportion of the cohort whose socio-economic situation is relatively limited and thus it links to primarily the level of income for the areas in which the students live, but might also act as an instrument for lower educational aspiration.

Educational aspiration was, as mentioned above, a mitigating factor for violence and crime. Thus if absences from school are high this would indicate that the most in danger cohort (those in secondary schools) for this forecasting window might face greater opportunity to be involved in crime and violent crimes in particular.

One issue to be dealt with is the different frequency compared to the daily/ minutely data associated with the various police systems and the various year commencing dates, financial, school or calendar.




Figure 3 Non-Working Households as a Percentage of the Area across the West Midlands

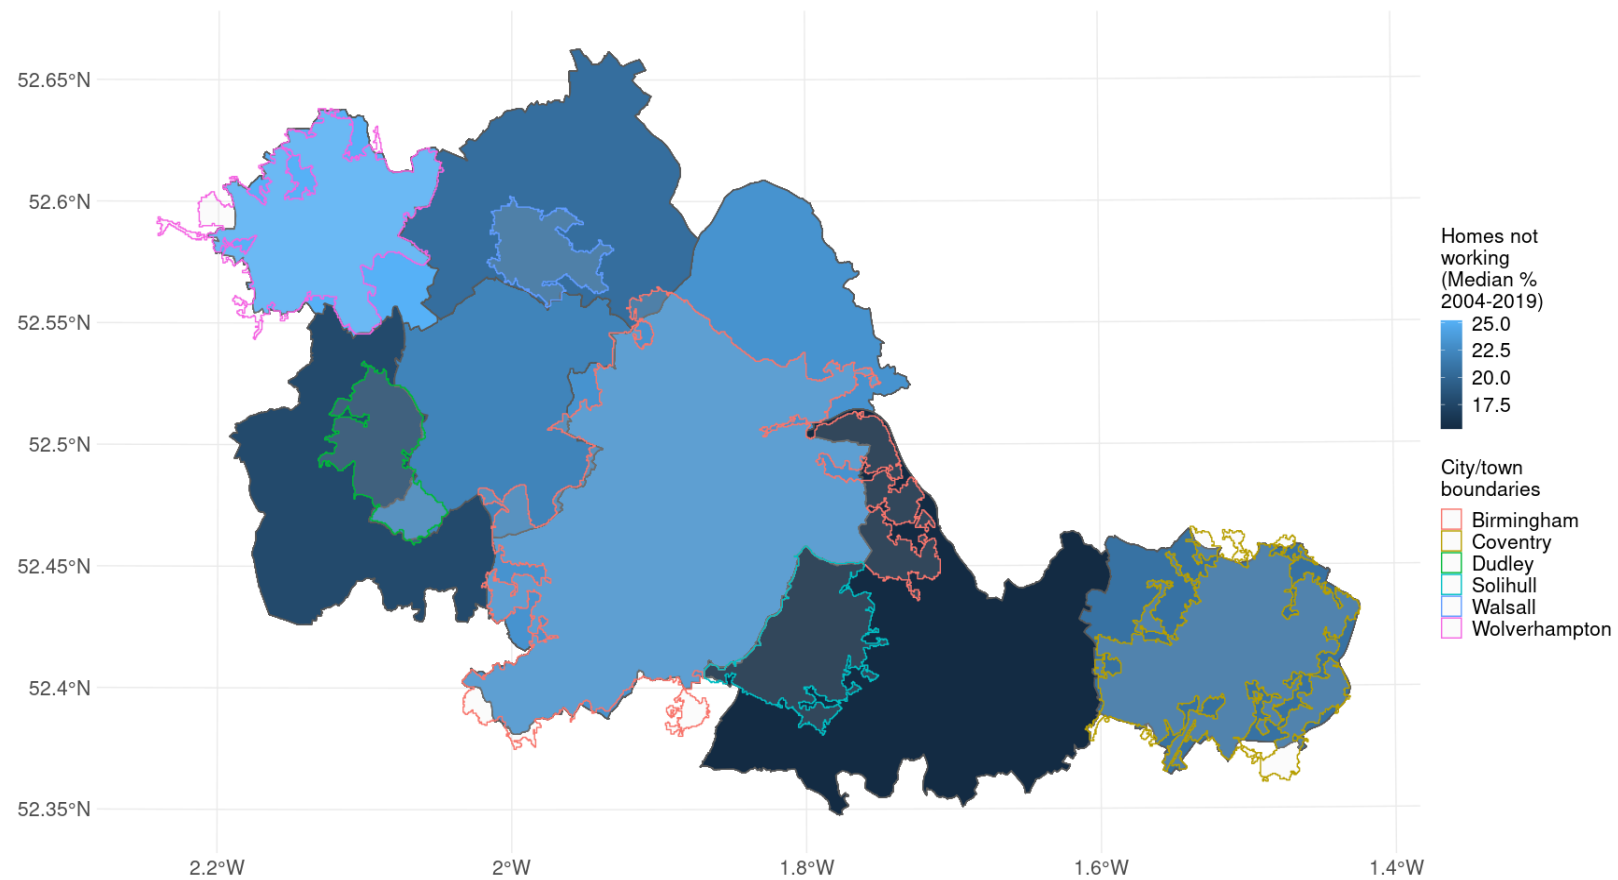


Figure 4 Local Authority Bounds As Defined in NOMIS and LGInform

When considering the impact of the various socio-economic factors, a transformation of the information is potentially useful. It might not be the absolute level of say, free school meals in an area, but how this level compares with the rest of the region. A mean normalisation is used. This allows us to see how the various towns are performing relative to the mean for the West Midlands. If the value is above 0, then the value is higher than the mean and if less than 0, it is lower.

$$x' = 100 \left(\frac{x_i - \bar{x}}{\max(x) - \min(x)} \right)$$

Free school meals were used to mirror the economic situation of the area. It reflects the economic situation of the parents. A child qualifies for free school meals if the parents are on income support, jobseeker's allowance, support as asylum seekers or various forms of tax credit. This is a measure of the impact of the area's economics on the children of the area. By considering the older children, one can assert that this is giving a measure of the expectations of outcomes by those students- seeing the family getting by and expecting to remain in a similar situation as they leave school. Higher levels of this variable is theorised to reflect a reduced opportunity cost of crime and potentially violence.

As with free school meals uptake, absence from school is an important measure of the educational aspirations and expectations of the area. This measure includes both authorised and unauthorised absences. Reasons such as illness are considered authorised, all other forms including turning up after the register is taken as unauthorised. This is measured as a percentage of half days missed within the local area. The same process as before (the mean normalisation) with the school meals used. Note that Birmingham's data is incomplete for the period considered and so the forecasts associated with this will be more uncertain. For interpretation, higher percentages of absences are indicative of more problems associated with the health and motivation of the students. This might be, *prima facie*, considered as a potential leading indicator of crime. Overall there has been a changing picture around the region; Wolverhampton has seen an improvement in the absences since 2014 as has Coventry (to a lesser extent). The towns in the sample are mostly clustered around the mean (as denoted by the dashed line) and are rarely the best or the worst in the region.




Figure 5 Free School Meal Qualification across the West Midlands




Figure 6 Standardised School Meal Measures across the Region




Figure 7 Absences from School across the West Midlands





Figure 8 Standardised Measure of School Absences

These variables are indicative of the socio-economic factors in the region particularly with respect to the younger elements of the population, who might move into crimes.

Linking the community to the local area and government, a measure of non-payment of council tax is also used. The reason for this is not specified; however the figure will include write-offs and arrears reflecting the areas' overall economic health.

Council tax not collected as a percentage of council tax due (2019/20) for All local authorities in West Midlands



Source:
Ministry of Housing, Communities & Local Government

Powered by LG Inform

Figure 9 Non Collection of Council Tax

The disparity across the region, with four of the five “worst” offenders being in the WMP area considered in this study is clear from Figure 9. It is these higher levels of non-payment that are used to reflect the potential issues in the WMP area.

One also needs to consider the political and broader local governmental incomes. These factors are picked up in two variables, non-payment of council tax and the police precept of the relevant areas. These reflect the legal and political aspects of the region and the ability and the desire of the council to pay for the police service.

It is not a clear that a simple political measure would suffice; there is no relationship between the various political parties at a more local level and the funding for law enforcement. Across the region there are a number of councils with different local requirements. These policing requirements are believed to be reflected in the precept granted to the force from each of those areas.

The Police Precept is not standardised by population; this means that Birmingham has a substantially higher precept but this is also reflected in the crimes levels. In light of the later techniques it is believed that this will not be a problem. Generally the precept is seen to be all but constant (allowing for inflation, which has been low over the period of interest ranging from 3.9%p.a. to 0.37% according to the World Bank and IMF (International Monetary Fund and files. (2020)).

In the various models considered, only three areas were explicitly used²; Wolverhampton, Birmingham and Coventry. These are generally the least well off and largest areas and given the ability for people to move around the region, these are used as explanatory variables.

² These areas were believed to be the most important urban areas and to be a focus of the West Midlands region.




Figure 10 Unpaid Council Taxes




Figure 11 Standardised Unpaid Council Taxes across the Region




Figure 12 Police Precept

5.2 Macroeconomic Data

To mis-quote Niels Bohr (as quoted in Ellis (1995),

Prediction is very difficult, especially when it is about the future (and economics is doubly fraught)

Niels Bohr & another

In order to have a sense of the economic outlook, data on the Government yields was included. The data was sourced from the Bank of England (2021). This is a monthly dataset based on the implied future yield. This is the rate or yield that is implied by the differences in spot rates now to the rate into the future. In this case, the forward rate implied by the yields now and those for 5 year Gilts. This is known as a yield curve (or would be if we had more than 2 points). The yield curve is often used as a good indicator of the probability of recession. Data collected from the Bank includes 5 and 10 year real implied yields. The real rate takes into account expected inflation. Thus the difference between the 5 and 10 year gilts yields will be suggestive of changing economic expectations. Using the longer yields than the forecast window allows us to consider the forward looking nature of budgeting by the various government agencies as well as the economic climate in 5-10 years. To consider the more contemporaneous information, local and central government interest rates are also available from the Bank. This gives a feeling for the ease of borrowing by the government at the current time³.

The data used looks to reflect aspects of the economic climate; we cannot predict the result of elections which is potentially overlapping in this timescale. Rather some idea of local & central government spending and borrowing is considered via the interest rates and requirements were possible. There are a number of parallel series here, potentially cointegrated (roughly speaking, they trend together), so these will be pruned to use only one in the model building.

Two lending rates were taken from this to represent the economic situation and ease of government fiscal expansion (to reflect the possibility of governments, local or national expanding spending). An effective date is created representing the date at which the interest rate is expected to be effective, so the five year rate is effective on a date five years after it was traded. This is directly used as the forecast of the interest rates facing the local area. On the other hand, the ten year rate is used contemporaneously to reflect a forward looking expectation about the economic situation more generally as in Taylor, Ratcliffe, and Perenzin (2015). This reflects the general outlook for the economy.

³ It is no longer feasible for local authorities to access the financial markets as they once did via the swaps market where they were ruled to have acted *ultra vires* (McKendrick (1997)).




Figure 13 Monetary Indicators

5.3 Crimes

Using the data from Crimes, those violent offences specified were selected from the initial SME request when the Lab considered serious violence. The current neighbourhood was used as a geographical grouping with a view to aggregating after an initial investigation.

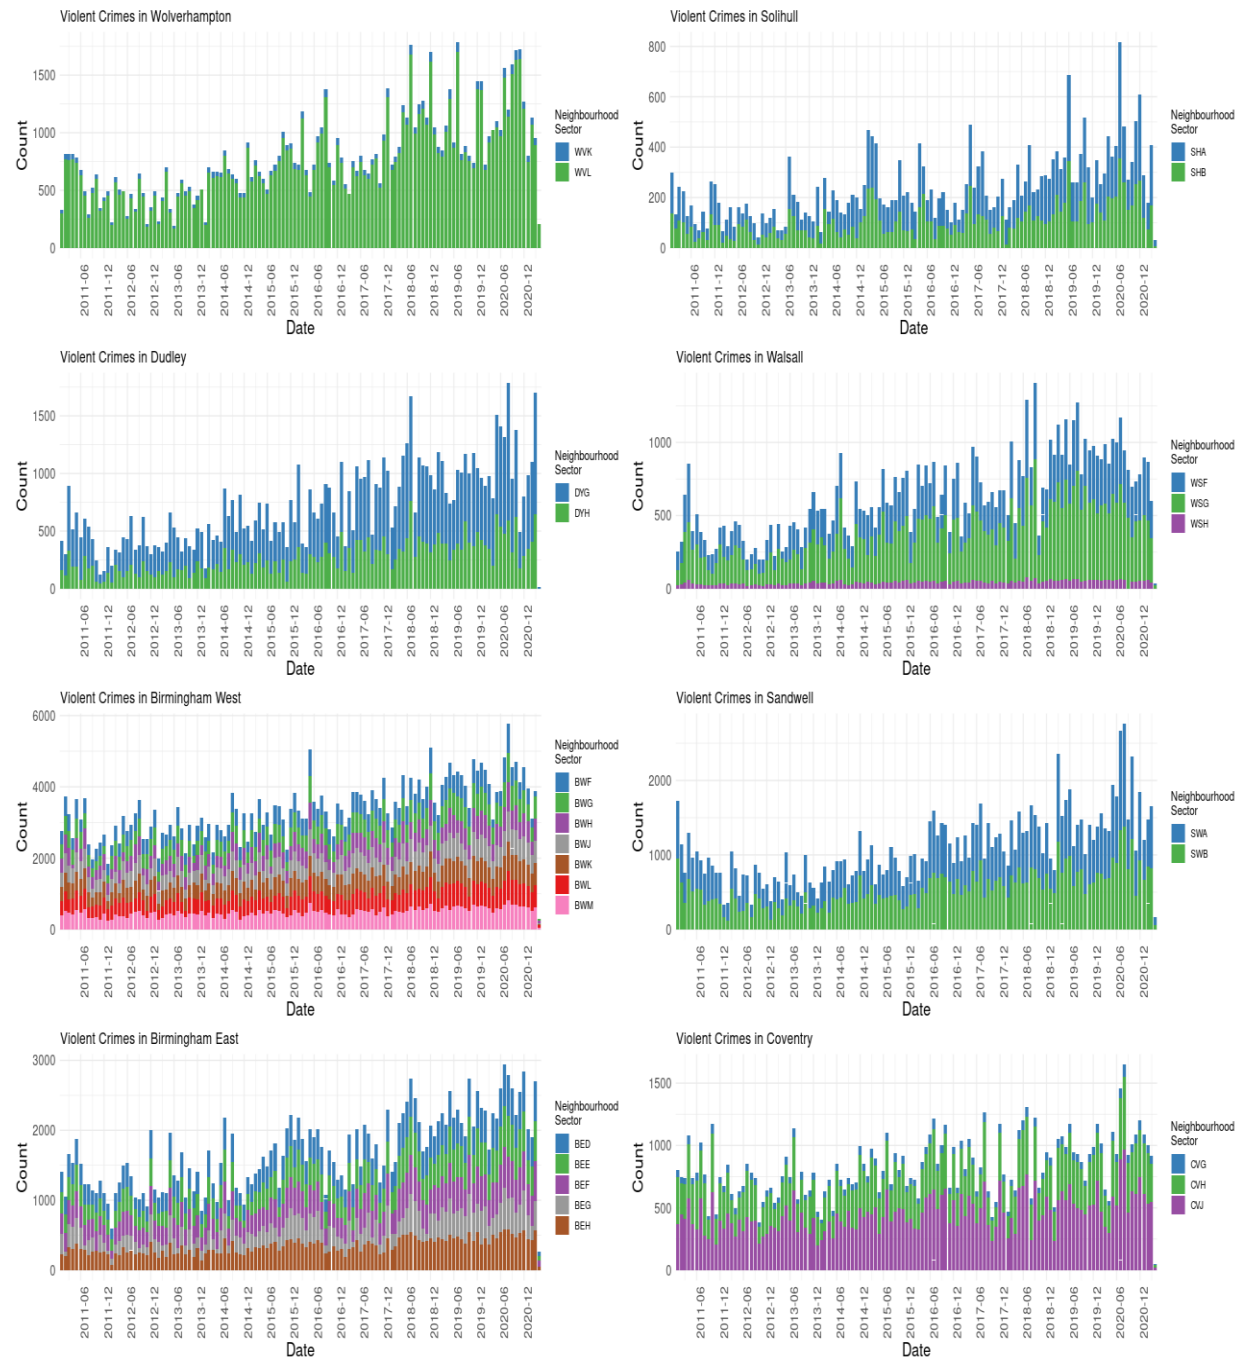


Figure 14 Violent Crimes in the Region

As would be expected with such a diverse area as the West Midlands, there is a variation across the areas. The three letters for the neighbourhoods are used as a brief indicator of the NPU split to see if there are any dominant areas. These are identical to the current

sectors. The LPUs also saw other, more general, crimes. A summary of these is presented below.

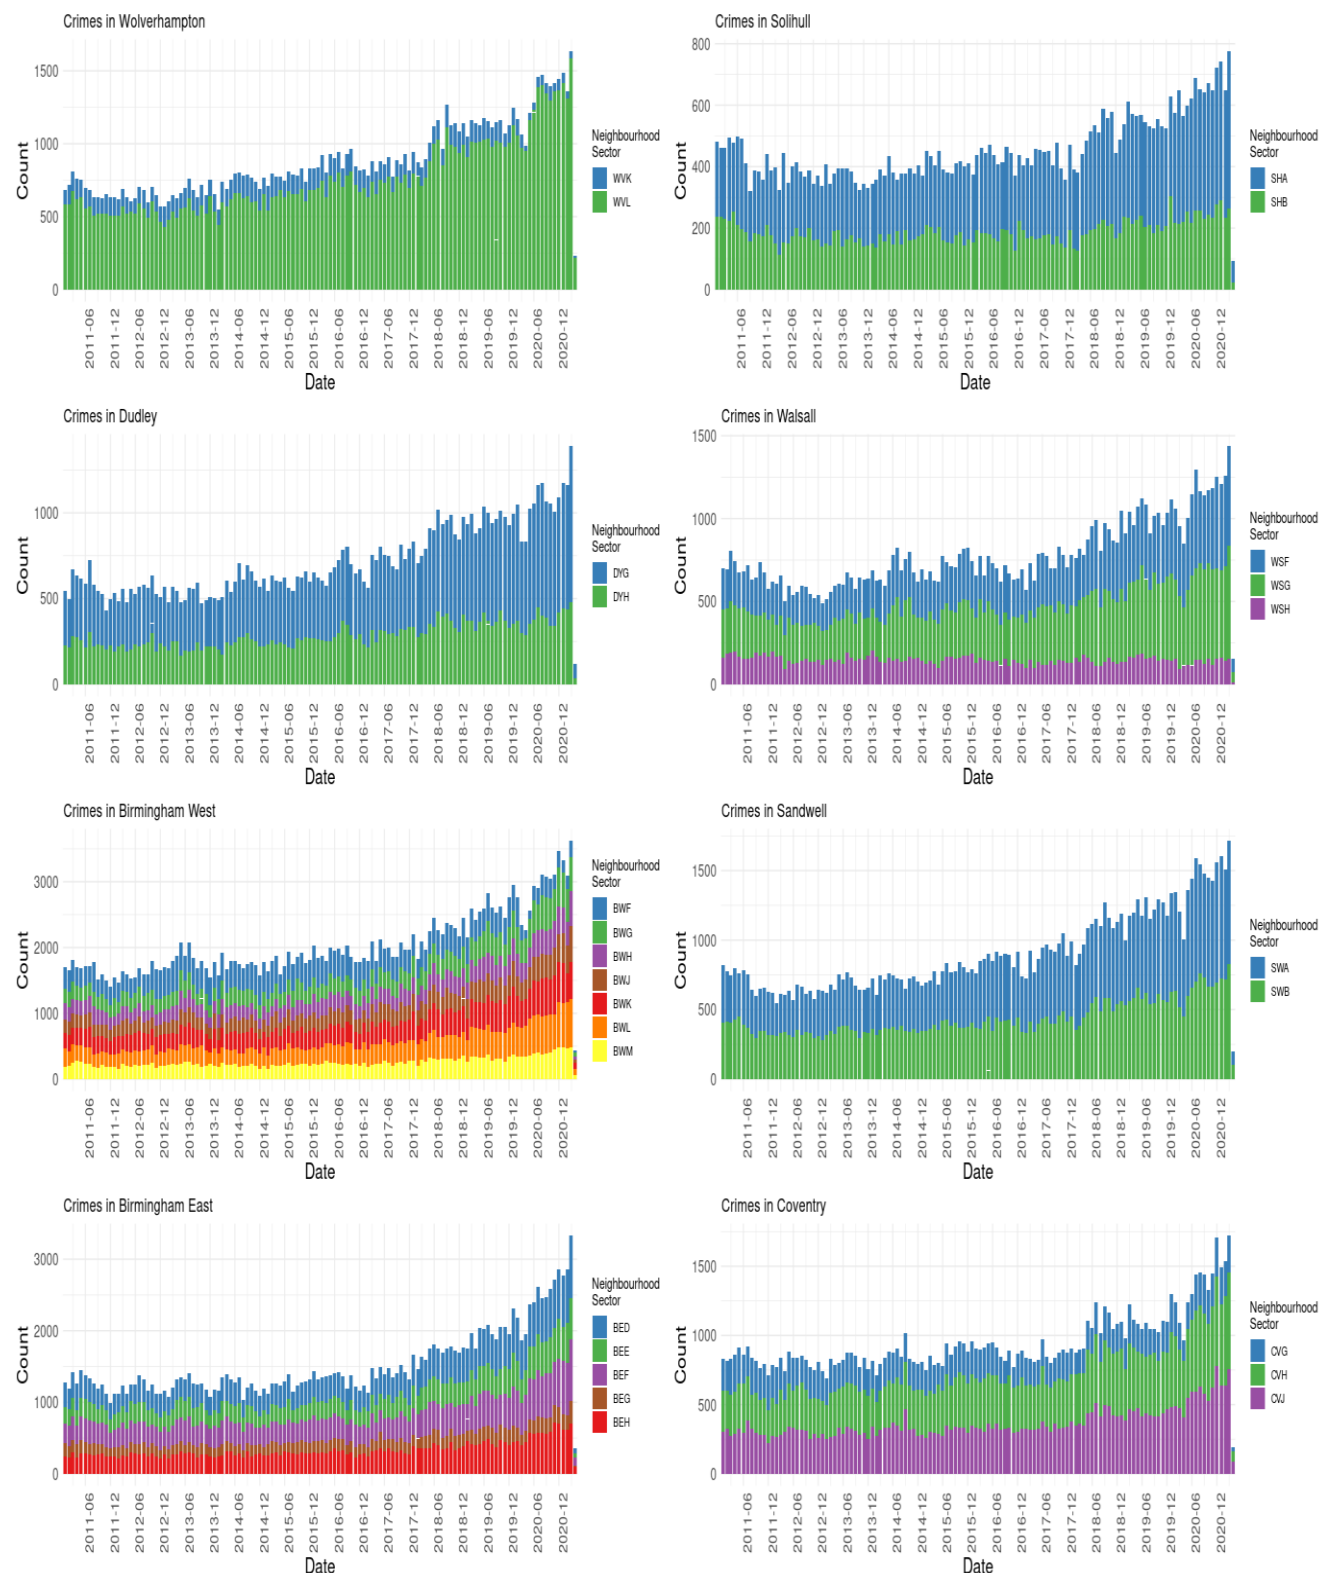


Figure 15 Crimes across the Region

6 Modelling

An initial modelling approach is to deal with the series individually. This allows us to consider the basic time series properties. Using an automatic/assisted ARIMA (Autoregressive Integrated Moving Average) algorithm based on the AICc will give the necessary unit root tests etc. before proceeding to the more complicated modelling approaches. Using multiple regressors rather than only one series within the ARIMA framework means that it is necessary to add forecasts for these variables, with the exception of the forward rate (which by its nature is forward looking). In order to do this, both an ARIMA and an ETS model (exponential smoothing state space model) were used to produce forecasts along with the 80 and 95% prediction intervals. These are to be fed into the main ARIMA/ARDL (Auto- regressive Distributed Lag) models for the variables of interest.

The data frequencies are different across a number of the series. In order to deal with these differences the annual values were spread across the monthly data. This does smooth out some of the seasonality potentially in the data but monthly data is otherwise not available. Unfortunately the usual MIDAS (Mixed Data Sampling) approaches tend to have the more frequent data as an explanatory variable and these are amalgamated into a single measure. This is not the case in this study; rather due to the nature of the relevant data it is assumed to be the same across all periods in question. This is perhaps most true of the monthly base rates for example but less so of the absentee-ism in schools. However this smoothing was considered a cost worth bearing.

This section looks at the main variable of interest, violent crime and looks at the effectiveness of building the univariate model and the multivariate model with explanatory features. The modelling used the corrected Akaike Information Criterion as the method of model selection. The use of the Schwarz Information Criterion made minimal practical difference, sometimes removing one lag, but not consistently.

6.1 ARIMA, ETS and Auto-Regressive Distributed Lag modelling

The univariate behaviour is an important factor in modelling. Autoregressive model gets its name from the Greek *αυτος* – referring to the self. Auto-regressive modelling is modelling using the data's own past values as the explanatory information. These models are one of the cornerstones of time series modelling and thus are an obvious place to start with the models. They include no other data than that of the series itself. This simplification will be built upon to include other factors and make the models multi-dimensional (such that one series will not be seen in isolation). Though it is a simplification, the story told by these models is useful; it gives an indication of whether the time-series are stationary (in distribution; stationarity roughly means that the probability of finding a particular value of a time series remains the same over the whole time period) and can advise the relevant lag structure of the models.

The model focuses upon the violent crime measures at a NPU level. As some of the data to be added is at best quarterly, and normally annual, monthly data is used in this element of the modelling. In order to look at the information in the data, NPUs are used as the main forecasting unit in the first place.

An ARIMA is defined based on the number of lags in the variable (p) and the error term (q) and the order of integration (d) which is the number of times the data needs to be differenced to create a stationary (in distribution) series. There are variants on this to allow for seasonal lags (a monthly series with an annual seasonal lag would have $t - 13$ involved). The model allows us to see what effect, if any yesterday's or last year's value had on today's. The choice of the number of lags in the relevant variables is based on the statistical information criteria and the model's test sample performance.

$$Y_t = \delta + \phi_1 Y_{t-1} + \phi_2 Y_{t-2} + \dots + \phi_p Y_{t-p} + \epsilon_t + \theta_1 \epsilon_{t-1} + \theta_2 \epsilon_{t-2} + \dots + \theta_q \epsilon_{t-q}$$

ARDL models are Auto-Regressive Distributed Lag (henceforth ARDL) models (Pesaran, Shin, and others (1995)). This is a modelling approach that allows both y_{t-i} and the lagged independent variables (x_{t-i}) to appear in the regression equation whether or not they are both integrated with order 1 or not. This property allows us to side-step issues of cointegration which are problematic with the relatively short series we have available. The ARDL model has a number of specifications, all of which basically say that the current value of the variable of interest (violent crime) is dependent upon past values of that variable *and* other variables up to and including the period which one is forecasting.

$$\begin{aligned} \phi(L)y_t &= \delta + \theta(L)x_{it} + v_t \quad \text{where L is the lag operator} \\ y_t &= \delta + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \dots + \theta_0 x_{t-0} + \theta_1 x_{t-1} + \dots + v_t \end{aligned}$$

One of the useful aspects of the ARIMA and ARDL families of models is that they give a relatively straight forward method of examining the influence of other factors. Other approaches that might be considered are the likes of Threshold models. In these if a variable passes a threshold, the process changes. This might be self-exciting (using the variable itself) or non-self-exciting, where another variable would be used as the threshold variable.

Though potentially useful, the requirement would be that all the variables would need to be tested across a number of lagged values and each could potentially be a threshold candidate. This would lead to an explosion of nuisance variables. With an aim of the modelling being parsimony, a decision was made to not investigate this approach. For a full description of TAR type models please see Tong (2011).

The other popular time series approach of Exponential Smoothing is a purely time series model. This approach splits the series into a number of different elements, trend and seasonal and models these. Each of these elements can take a number of different forms, hence there are 9 different specifications of which the Holt-Winters (Holt (2004) and Winters (1960)) is perhaps the most well known in economic circles. This model is a weighted average of the last occurrence and the forecast.

Simplest

$$\hat{y}_{t+1} = \ell_t = \alpha y_t + (1 - \alpha)\ell_{t-1}$$

Adding a Trend to the Smoother

$$\hat{y}_{t+1} = \ell_t + hb_t$$

Levels

$$\ell_t = \alpha y_t + (1 - \alpha)(\ell_{t-1} + b_{t-1})$$

Trend

$$b_t = \beta^*(\ell_t - \ell_{t-1}) + (1 - \beta^*)b_{t-1}$$

These models smooth the data depending on the parameters, α and β . The lower α and β^* are, the smoother the averaging becomes. These coefficients are estimated using a likelihood approach under a normal distribution. The models are selected using information criteria in the same manner as the ARIMA style models.

6.2 Multivariate Models

These models are univariate time series (excepting ARDL), where each series is seen in isolation. In economics in particular but in other fields too, a multivariate approach can be beneficial as this allows both series to effect each other allowing for example Birmingham East to impact Birmingham West. These models were developed by (Sims (1980) in an atheoretical framework which suits forecasting the form of series here where there is no particular theoretical framework in which to work.

These are normally constrained to the variables of direct interest, though adding exogenous variables is possible. Of course the same issues of requiring certain data to be pre-forecasted remain. The standard Vector Autoregression requires stationarity, however the ARDL extension is a relatively natural approach to allow mixed orders of integration to co-exist.

This is often written in *vector* notation, but it is helpful to see the underlying structure in a smaller case. Other forms such as the structural VAR are also possible but these are equivalent.

$$\begin{aligned} \begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} &= \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} + \begin{bmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{bmatrix} \begin{bmatrix} y_{1,t-1} \\ y_{2,t-2} \end{bmatrix} + \begin{bmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{bmatrix} \\ \Leftrightarrow Y_t &= AY_{t-1} + \epsilon_t \end{aligned}$$

This demonstrates the inter-relationship between the variables. It is also possible to write these more complex functions in terms of a single lagged term, which makes inference more straightforward. Though Granger Causality⁴ can be examined in this framework, this work will *not* consider it.

A further advantage of using the VAR framework is the impulse response function (IRF). This shows the impact of one of the variables changing across the other variables; for this project this would take the impact of a shock to violent crimes in Birmingham West on all the other areas. Further this allows us to evaluate the impact over time; how long it takes for a shock to no longer have an effect across the various areas.

6.3 Cross-Validation

It is not possible to use the standard cross-validation techniques for selecting a model on time series data. To do so would destroy the time series properties of the data; thus instead of selecting sub-groups of the data, a rolling forecasting origin is used. There is

⁴ If x precedes y, then x Granger causes y. The qualification of “Granger” causes is due to the possibility of other underlying (hidden) causes that, for example, cause x and therefore the full chain of causality is not necessarily known. The classic example is that Christmas cards Granger cause Christmas, but do not cause Christmas with the latent variable of Christ’s birth not appearing in the causality link.

in effect a window of increasing size from which the forecasts are made. The data is stretched to give the new data with the starting point fixed at the starting point for the first dataset (rather than a rolling window of say 10 months). The results of this approach will give rise to the relevant model for the forecasting.

7 Results & Analysis

The approach discussed above was used to forecast violent crimes across the areas of the Service. It begins with a brief consideration of the time series properties, moving on to the univariate time series. Finally it deals with the data in a multivariate setting, both only using the dependent variables (violence levels) and introducing the other explanatory levels. Results are presented in terms of standard forecasting measures where applicable, with other information that might be of use also presented as appropriate.

7.1 Univariate Analysis

This section considers the time series in isolation, whereupon no one area has any impact on any other. This is clearly a simplification but it is a good place to start the process to understand the nature of what is occurring. The first step is to look at the temporal relationship of the data, how one month's data effects the next. This is demonstrated through the use of auto-correlation functions (ACF) or graphs. A line that remains high, near one, suggests that even the distant past will continue to have an effect on today. A case such as this is extreme. A steady decline is more common. In this form, the autocorrelation does not strip out the previous months' impacts. This task is performed by the partial auto-correlation function (PACF). For the main time series, these functions are presented below.

As a rule of thumb, if the autocorrelation function tails off and the partial is more precipitous then the model is likely to be an auto-regressive model rather than a moving average model. However this is also the case in integrated time series which show substantial hysteresis or memory⁵. In order to consider these problems (statistically speaking), the differenced data is also investigated and found to be stationary. Further examining these graphs there is little evidence of any seasonality, with the various estimates of 6, 12 and 18 months all low. This would suggest that there is a degree of statistical non-stationarity in the data, which means that looking at the changes will be helpful.

⁵ The KPSS test (Kwiatkowski et al. (1992)) is preferred over the Dickey Fuller type test (Dickey and Fuller (1979)) as the null under the KPSS is stationarity of the series with an alternative of non-stationarity.

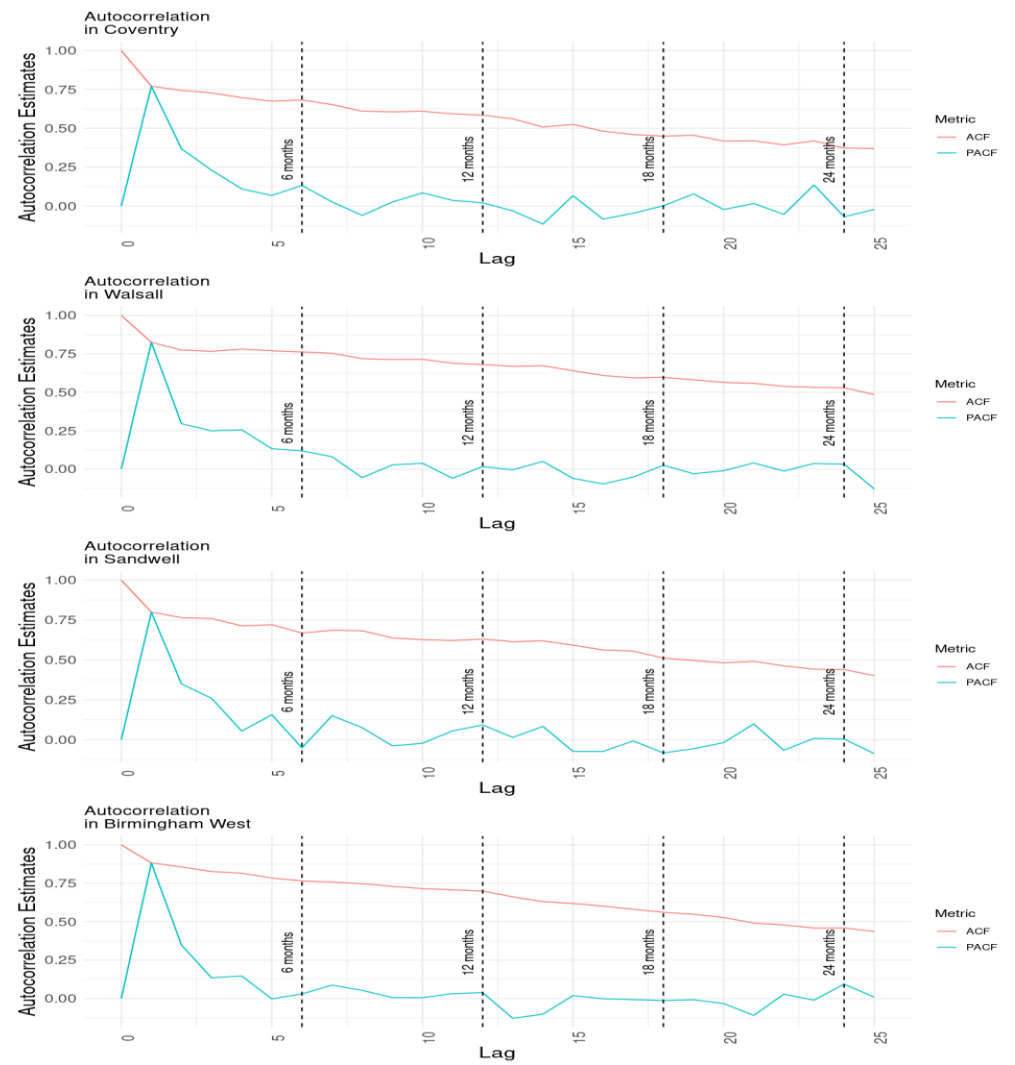
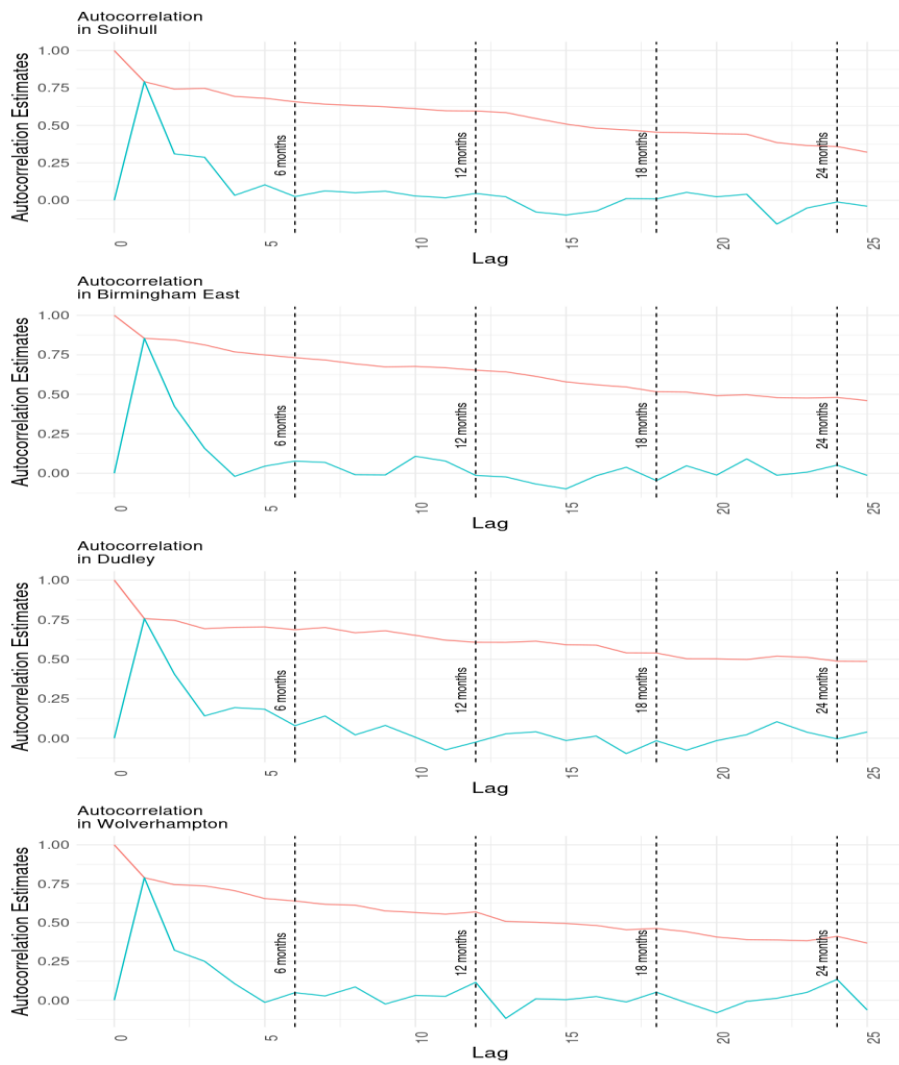


Figure 16 Time Series Properties of Underlying Levels Data

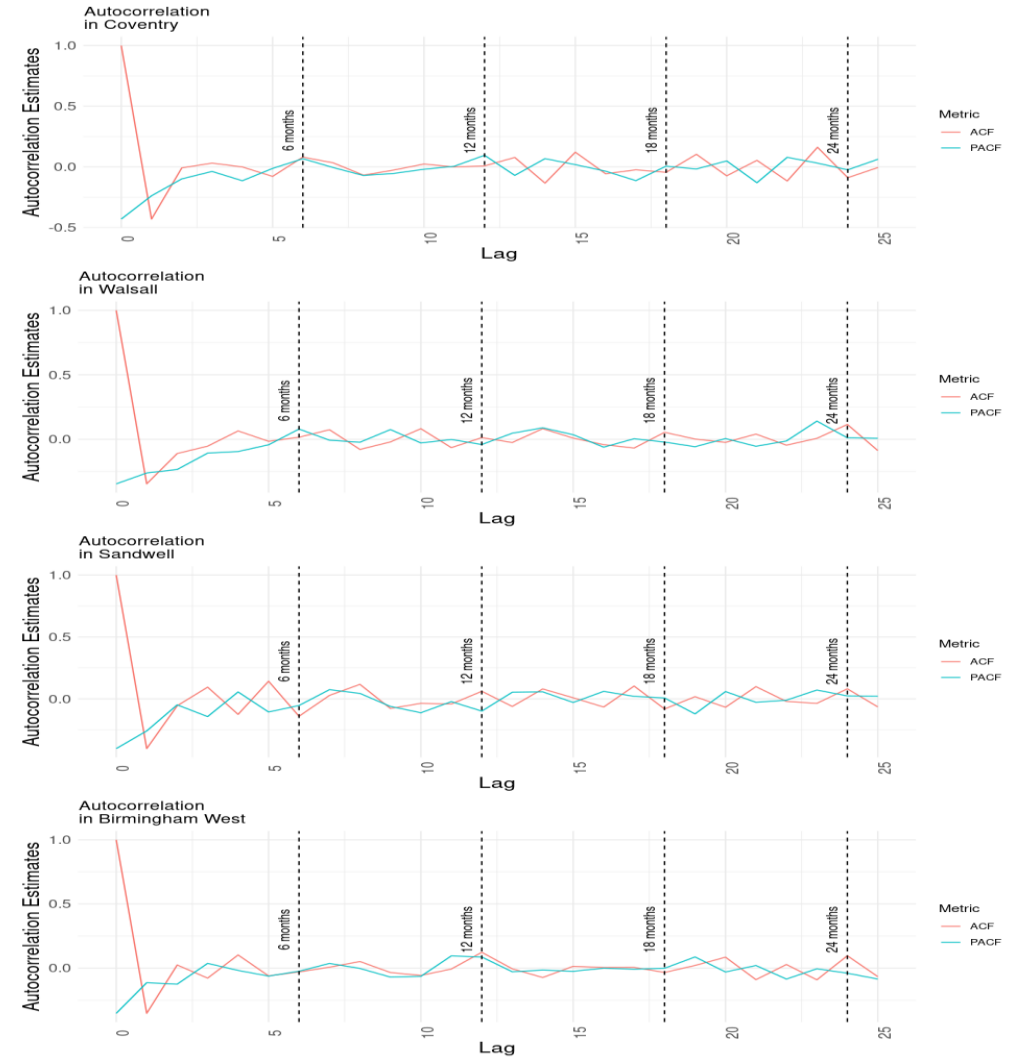
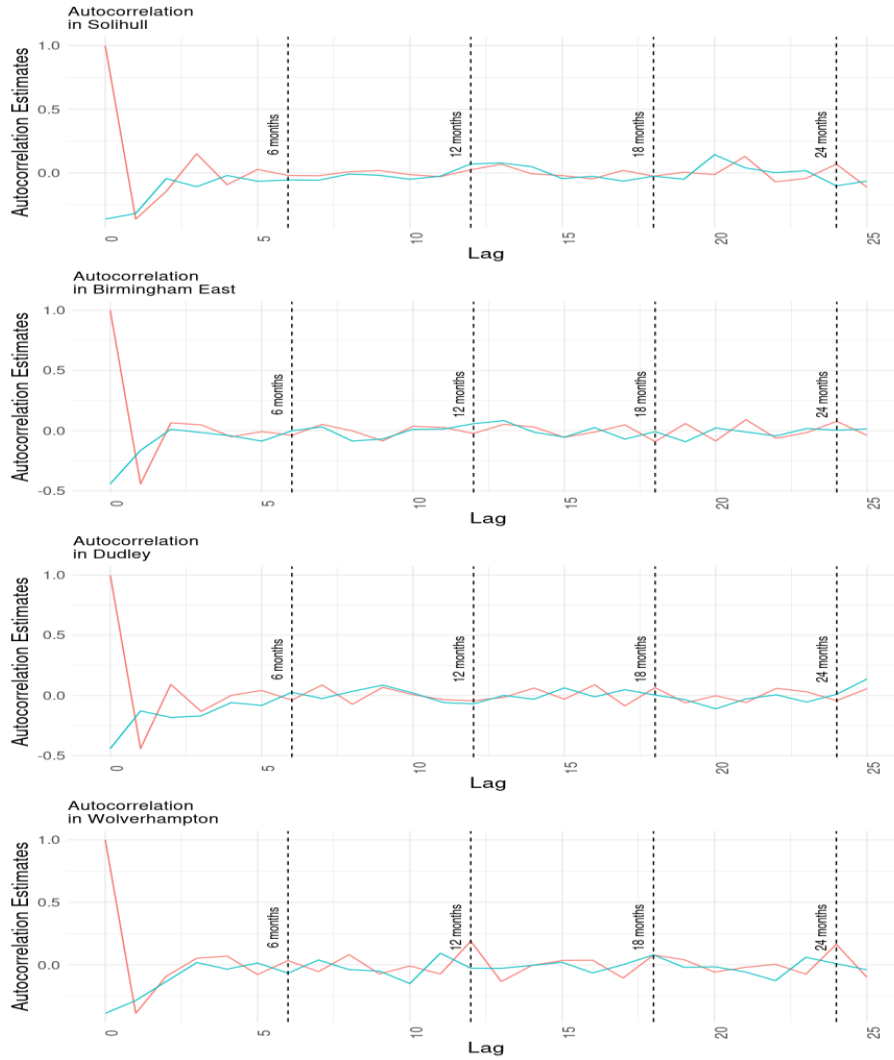


Figure 17 Time Series Properties of Underlying Differenced Data

As part of the investigation of the properties of the underlying, main series it is helpful to consider the (de-)composition of the series into seasonal and trend elements. These allow us to discover any major shifts in the underlying data and to consider any obvious seasonality based on an additive decomposition.

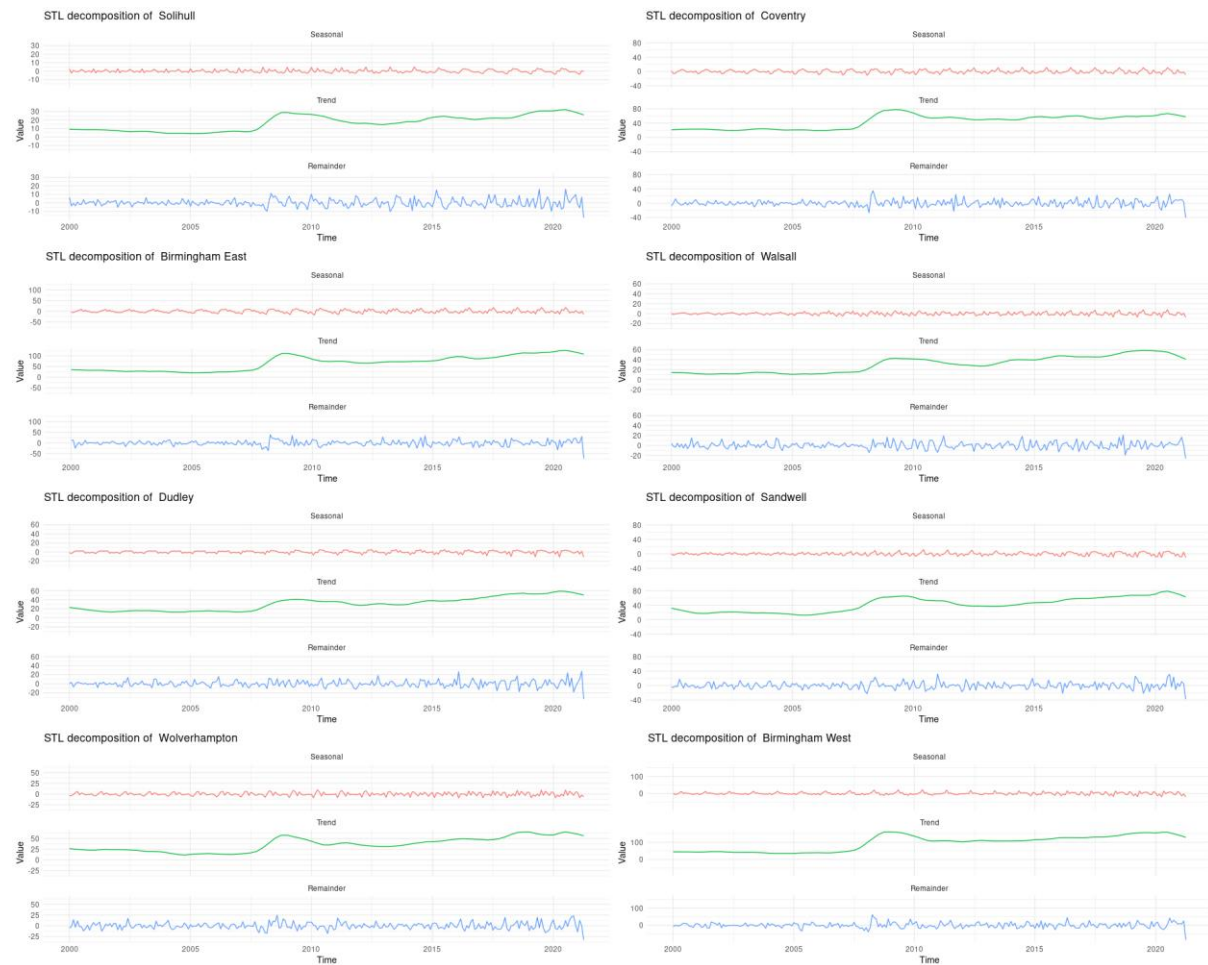


Figure 18 Decomposition of Underlying Data

7.1.1 ARIMA and Exponential Smoothing

The univariate ARIMA models are estimated in the usual manner by area. The data used for the modelling covers the period January 2000- April 2021 even though we can see that there is a split in the data in 2008. This will lead to some issues with the model.

The model coefficients are presented in the appendix. The structural break is a problem in that it can create a spurious unit root, though we can see where the break is this can lead to problems with the ARIMA type models.⁶ This is informative and can be mitigated by the use of the exogenous variables in the modelling.

The exponential smoothing model discussed elsewhere is also estimated and the mean error and other associated measures presented in Tables 1 & 2. The results are presented for the optimal model, given the data used and the information criterion used in the selection process.

The accuracy of the models are presented below. This uses the cross validation techniques in Section 6.3, where a rolling and expanding window is used with an average used across the windows based on both 1 and 5 year time frames. As one would expect, there is a decrease in accuracy across models on the increase in forecast window. The absolute values for the mean errors and percentage errors are large suggesting that there is a lot of noise in the forecasts and the models are likely to be unstable.

Table 1: Accuracy of 1 Year Forecasts

NPU	Model	ME	RMSE	MAE	MPE	MAPE	MASE	RMSSE	ACF1	Winkler	Percentile	CRPS
BE	ETS	2.176	23.026	16.579	-6.812	22.179	0.891	0.866	0.173	158.219	12.968	12.849
	ARIMA	1.932	23.709	17.502	-7.272	23.184	0.941	0.892	0.227	131.874	12.723	12.605
BW	ETS	2.402	28.109	19.417	-8.592	22.236	0.878	0.855	0.273	175.954	14.565	14.433
	ARIMA	1.921	27.647	19.332	-8.566	21.823	0.874	0.841	0.269	171.040	14.495	14.361
CV	ETS	0.223	16.471	12.284	-15.48	33.156	0.948	0.893	0.236	99.544	9.357	9.266
	ARIMA	0.796	15.861	12.034	-14.03	31.779	0.929	0.860	0.266	85.914	8.764	8.680
DY	ETS	2.485	13.523	10.227	-15.18	35.927	1.014	1.062	-0.02	83.597	7.486	7.417
	ARIMA	2.454	13.526	10.229	-15.26	35.947	1.014	1.062	-0.02	83.592	7.487	7.418
SH	ETS	0.807	8.234	6.102	-6.471	26.852	1.017	0.971	0.311	49.525	4.543	4.499
	ARIMA	0.631	8.126	5.999	-6.923	26.219	1.000	0.958	0.255	48.960	4.459	4.417
SW	ETS	2.810	17.414	12.451	-4.690	23.801	1.010	1.053	0.274	120.047	9.389	9.303
	ARIMA	3.071	17.213	12.342	-4.108	23.445	1.001	1.041	0.232	117.920	9.200	9.117
WS	ETS	0.299	11.437	7.986	-12.36	25.868	0.844	0.905	0.174	73.485	5.985	5.930
	ARIMA	-0.12	11.723	8.271	-13.33	26.616	0.875	0.928	0.194	75.501	6.143	6.087

⁶ Much of the analysis of the impact of structural breaks is based on Dickey Fuller style tests, with the null of non-stationarity rather than under the null of stationarity as in KPSS.

NPU	Model	ME	RMSE	MAE	MPE	MAPE	MASE	RMSSE	ACF1	Winkler	Percentile	CRPS
WV	ETS	1.540	16.060	12.247	-5.157	24.027	1.046	1.001	0.367	93.065	9.036	8.950
	ARIMA	0.961	16.617	12.884	-6.139	25.109	1.100	1.036	0.388	92.326	9.449	9.358

Table 2: Accuracy of 5 Year Forecasts

NPU	Model	ME	RMSE	MAE	MPE	MAPE	MASE	RMSSE	ACF1	Winkler	Percentile	CRPS
BE	ETS	12.154	29.960	23.365	-0.656	29.020	1.256	1.127	0.272	227.288	18.948	18.766
	ARIMA	12.526	30.083	23.644	-0.133	28.846	1.271	1.131	0.294	161.640	16.825	16.663
BW	ETS	12.880	35.472	26.307	-6.036	30.024	1.190	1.078	0.329	212.650	19.450	19.265
	ARIMA	12.410	35.259	26.165	-6.192	29.793	1.183	1.072	0.335	211.995	19.442	19.256
CV	ETS	1.040	22.771	15.481	-19.93	43.785	1.195	1.235	0.535	195.546	15.314	15.166
	ARIMA	2.706	17.013	13.016	-15.55	37.184	1.004	0.923	0.270	95.316	9.693	9.599
DY	ETS	7.498	17.160	13.297	-12.74	46.054	1.318	1.348	0.040	94.228	9.621	9.529
	ARIMA	7.456	17.141	13.279	-12.83	46.039	1.316	1.346	0.040	94.026	9.612	9.520
SH	ETS	4.252	10.228	7.881	4.389	30.796	1.313	1.206	0.340	54.275	5.682	5.627
	ARIMA	4.130	10.162	7.825	3.953	30.571	1.304	1.198	0.326	53.745	5.623	5.568
SW	ETS	7.830	20.937	15.326	0.130	27.880	1.243	1.266	0.287	125.839	11.608	11.496
	ARIMA	8.910	21.339	15.723	1.868	28.156	1.275	1.290	0.276	121.256	11.567	11.456
WS	ETS	3.703	13.294	9.935	-10.90	33.420	1.051	1.052	0.299	76.793	7.222	7.153
	ARIMA	3.424	13.523	10.139	-11.75	34.140	1.072	1.071	0.323	78.274	7.344	7.275
WV	ETS	6.602	18.845	14.341	1.601	27.052	1.225	1.175	0.394	92.827	10.621	10.517
	ARIMA	5.716	19.010	14.682	0.343	27.493	1.254	1.185	0.413	102.887	10.999	10.892

The ARIMA model specifications from the cross validation were used to model the overall series as were those of the exogenous factors where applicable; where ETS is used this is taken to be the $t, t - 1, \dots$ values for models keeping the ARIMA based orders where these forecasted values are needed. The coefficients of the cross-validation models for the ARIMA models are provided in the appendix.

There is stability across the coefficients and structures. The order of the ARIMAs is used in the final estimation of both ARIMAs and ARDLs for the next steps. The maximum value of these gives the approximate order required for the upper end of the search on the estimation with those where the measures coincide setting the order of that parameter. All these ARIMA models are giving a first difference, so the models are not directly of the levels but rather they consider the change in the value and they mostly use the moving average element rather than the auto-regressive approach. There is some evidence for seasonal moving averages as well, though rather less than the main equation.

Table 3: Model Orders Generated through Cross- Validation

NPU	Mean AR lag	Min AR lag	Max AR lag	Order of Differencing	Mean MA lag	Min MA lag	Max MA lag	Max Seasonal AR lag	Seasonal Differencing	Mean Seasonal MA lag	Max Seasonal MA lag
BE	0	0	0	1	1.0909	1	2	0	0	0.6970	1
BW	0	0	0	1	1	1	1	0	0	0.742	1
CV	0	0	0	1	1	1	1	0	0	0.015	1
DY	0	0	0	1	1.045	1	2	0	0	0	0
SH	1.32	1	2	1	1.682	1	2	0	0	0.864	1
SW	0.33	0	1	1	1.909	1	2	0	0	0.803	1
WS	0.44	0	2	1	2.030	1	3	0	0	0	0
WV	1.818	0	3	1	0.967	0	1	1	0	0	0

From these parameter estimates, a final VAR and an ARDL model was estimated using the data from the other data sets. Note the mean orders not being an integer is not a problem. It signifies that there was some variation across the estimates and the deviation of these from the maximum gives an indication of the spread.

7.2 VAR and ARDL Models

This section considers the results of the forecasting and estimation of the multivariate relationships, where each area effects each other area. The VAR is presented first. This only includes the dependent variable and its lagged values; the ARDL (or more correctly VARDL) is the final set of results.

7.2.1 Vector Auto Regressions

Cross validation of the Vector models were also performed. As before this gives the behaviour of the coefficients and the accuracy associated with the estimation over the sample. This information informs the general VAR fit and that of the ARDL models. It should be noted that due to the short time scale of the local authority type data cross-validation was considered not sufficiently effective; when there are only 10 observations estimating on 7 and forecasting a rolling window will be more inaccurate than basing the model on all 10 observations.

The VAR accuracy measures using the 1 year and 5 year windows are presented below. As with the univariate models these suggest an order of Autoregression. The impulse responses for each of the equations are also presented. These IRFs are reflective of the median response across a number of different orders of the model specification.

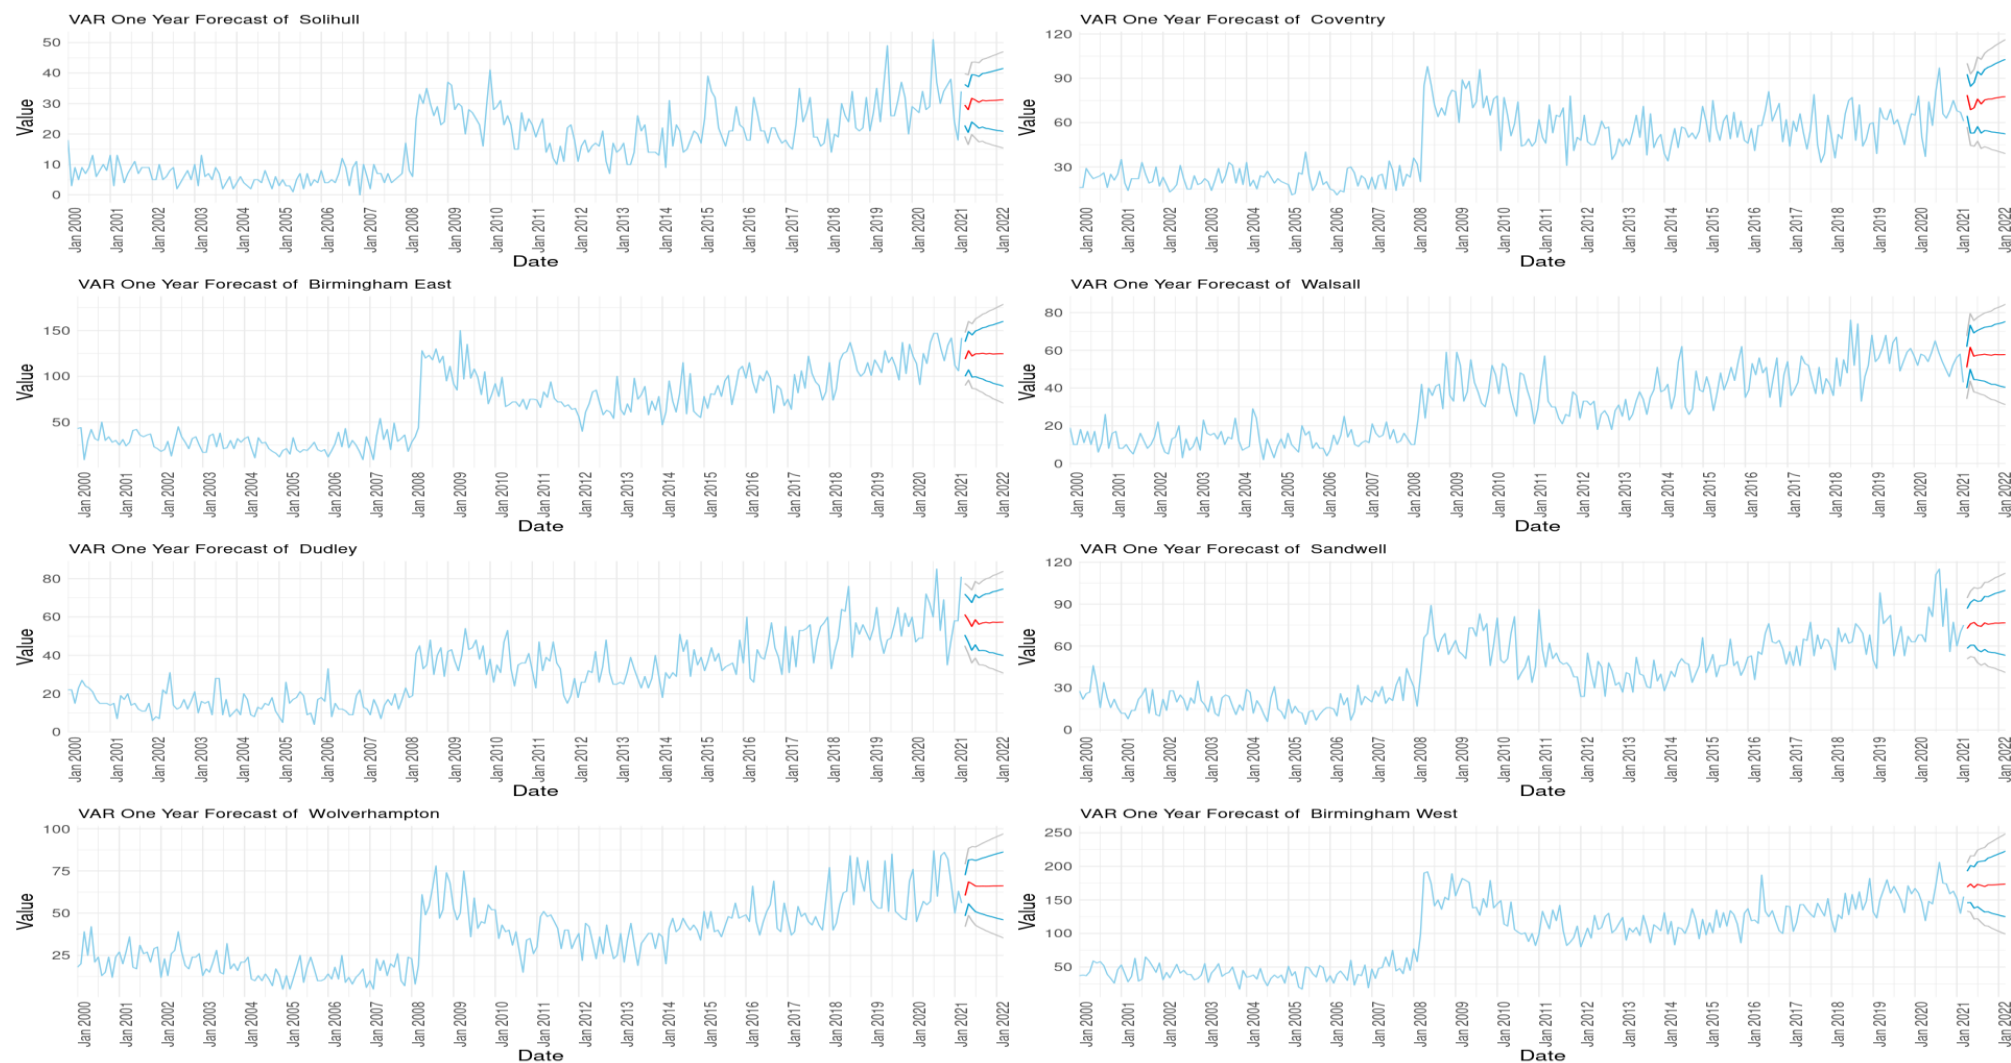


Figure 19 One Year Forecasts for VAR model

Key: red mean forecast, sky-blue 80th % interval, grey 95th % interval

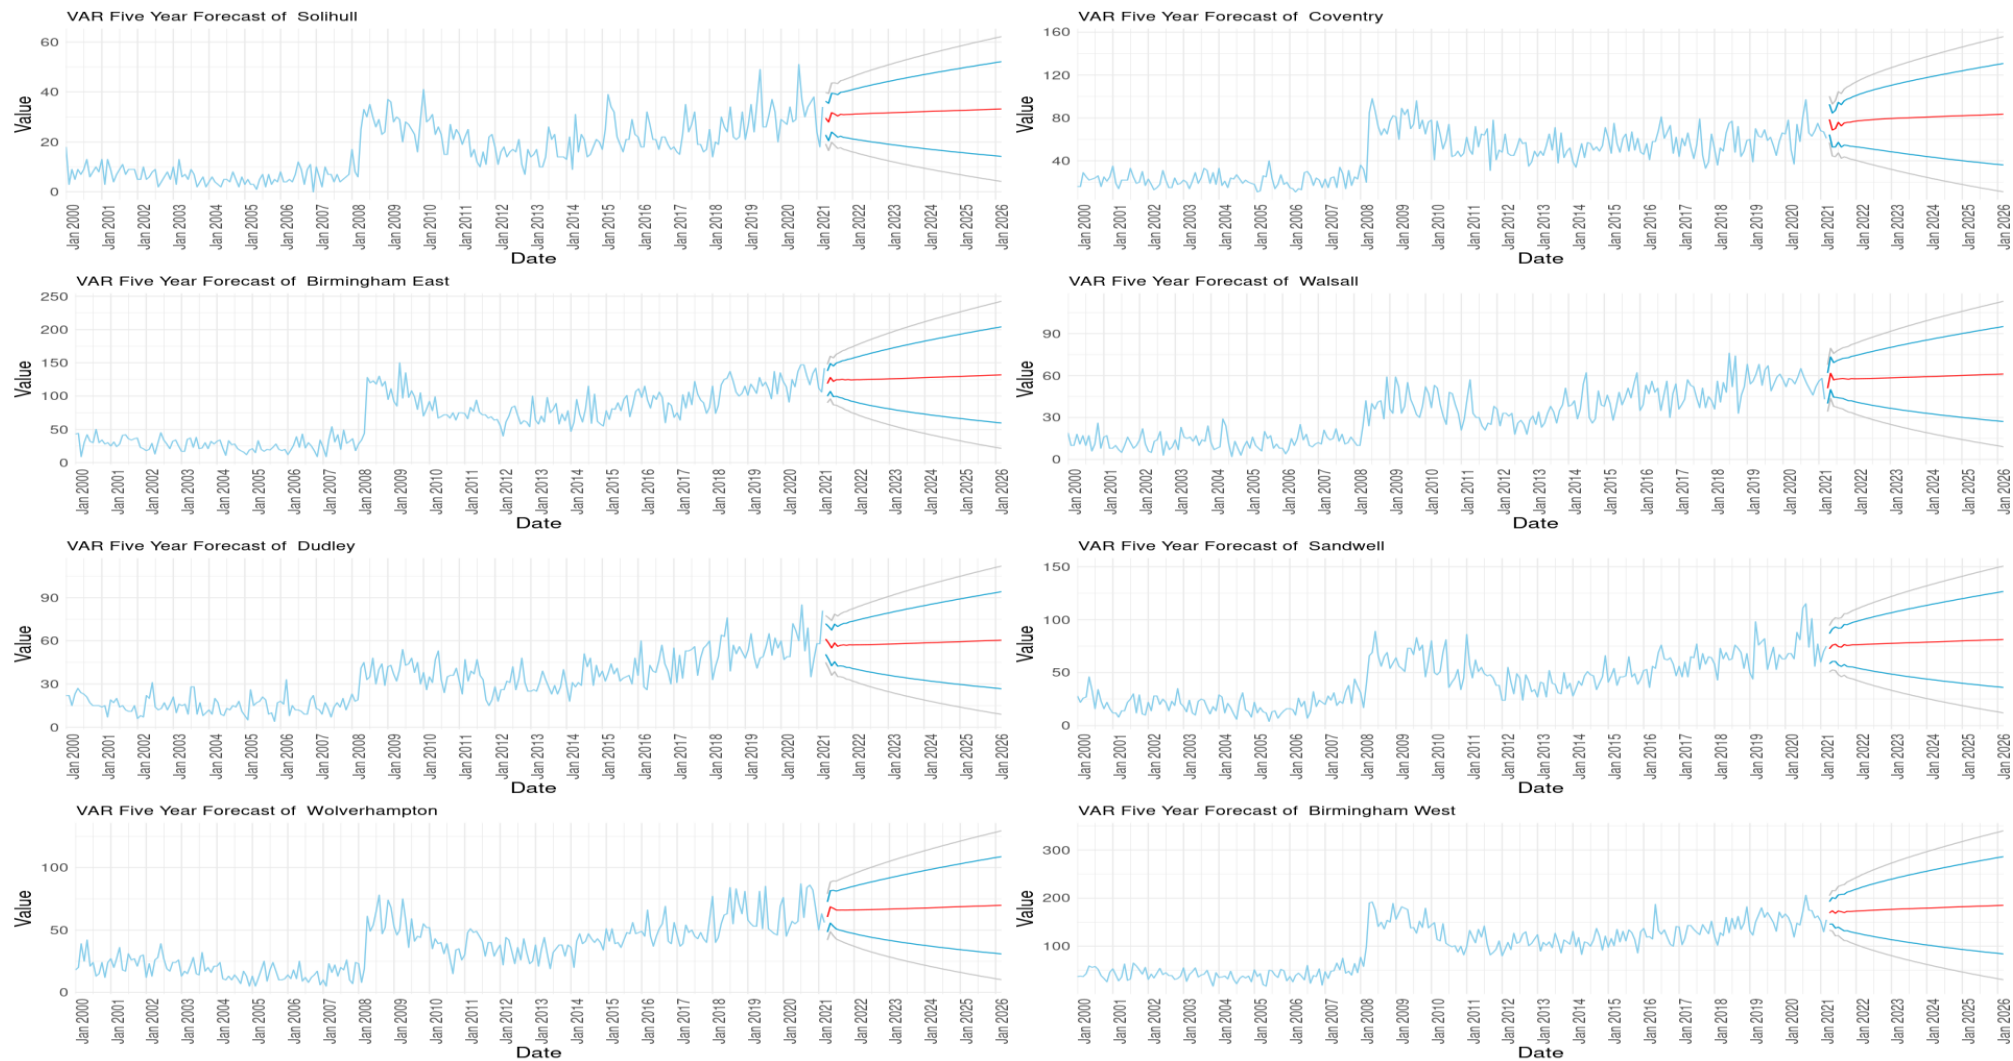


Figure 20 Five Year Forecasts for VAR Model

Key: red mean forecast, sky-blue 80th % interval, grey 95th % interval

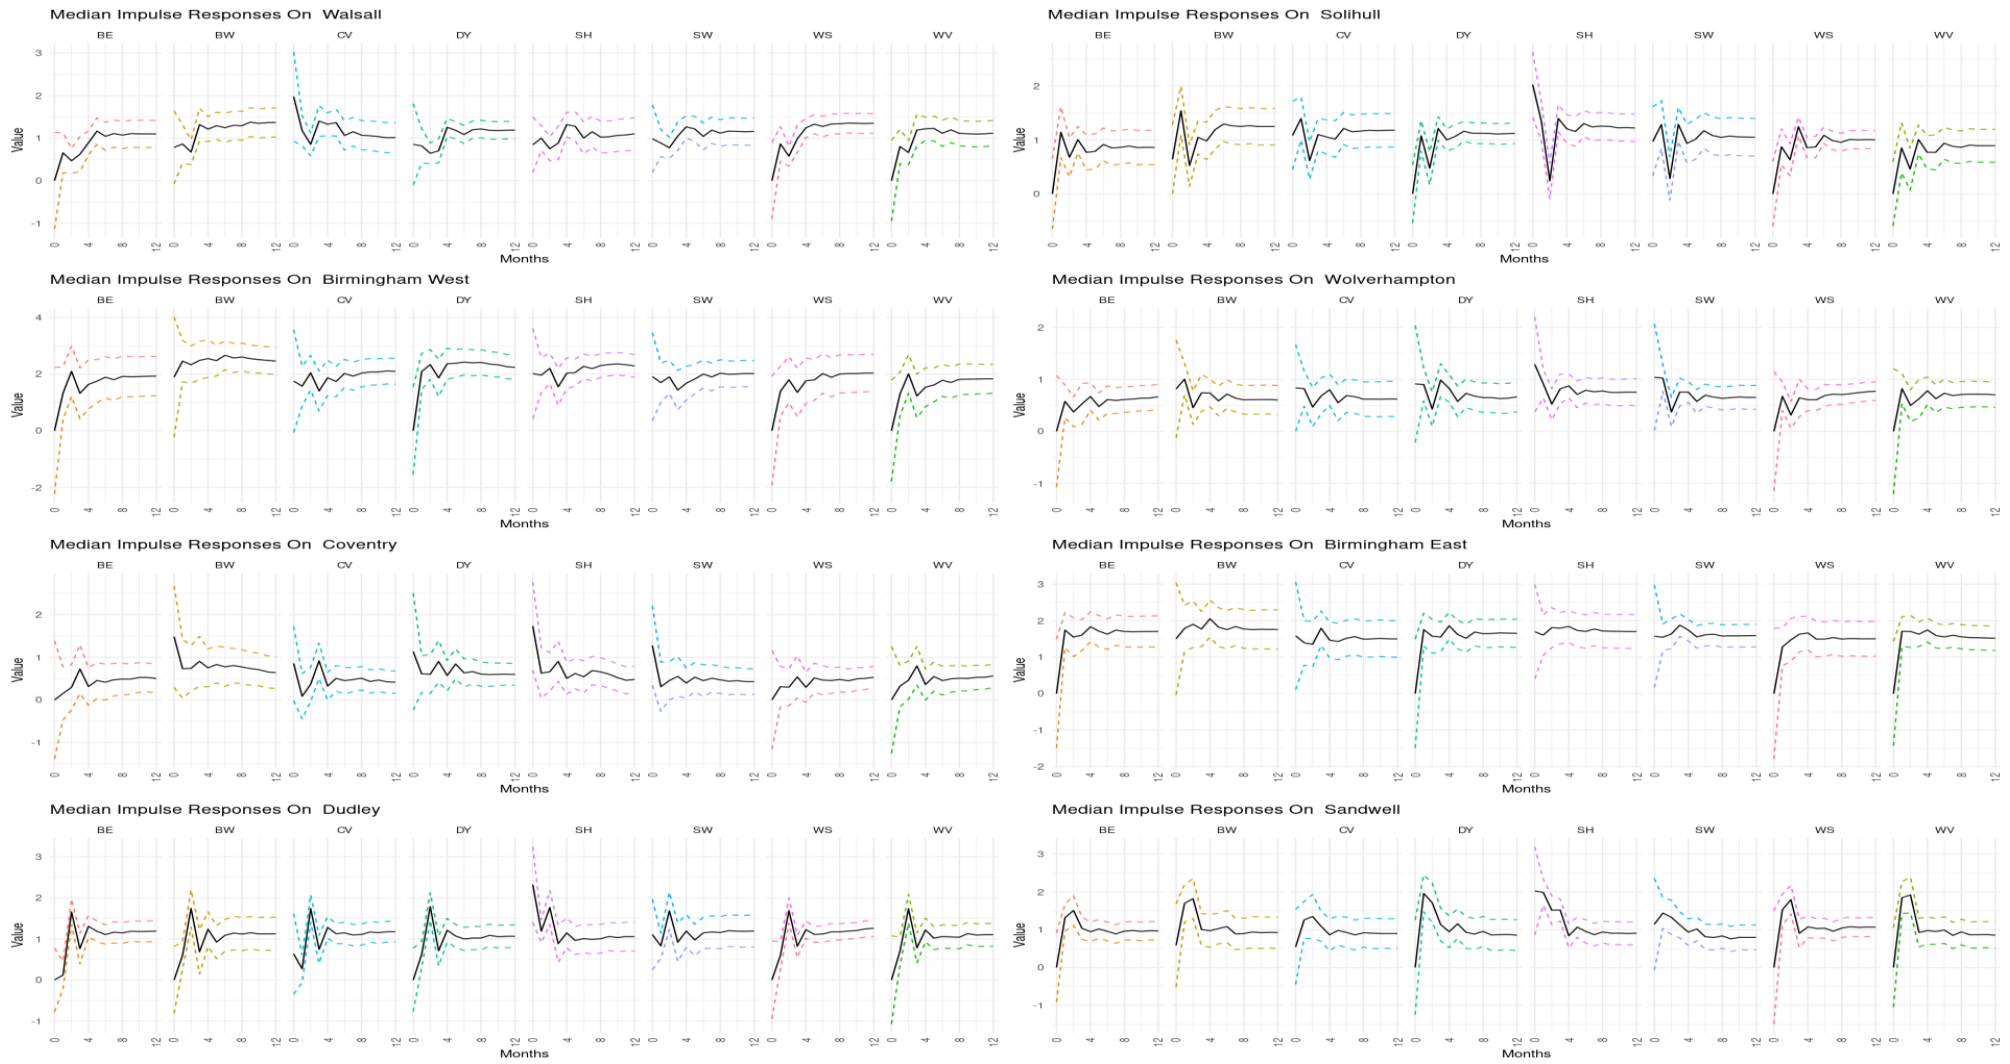


Figure 21 Impulse Response Functions for VAR models

The initial VAR is likely to be not well specified as it is based on a non-differenced series, however it is beneficial to consider it, not least to check that the findings of the previous stages are still valid. Further it does highlight some of the cross-NPU impacts. There are a very large number of coefficients in these models and so they are presented separately (in the appendix). The final observation (April 2021) was removed as this was an incomplete month. Removing this removes any downward bias based on this measurement period.

As we would expect the forecast intervals are large and there is in essence a straight line projecting from the end of the series. The VAR in levels is however problematic in that some structure still exists in the residuals⁷. We can see this in the autocorrelation charts as presented for other series. Once this is discovered, and given the information from the univariate series, a single differencing was used.




Figure 22 Autocorrelation of VAR model

A final concept worth looking at is the Impulse Response Function. This is the effect over time on the outcomes; it answers the question when crime in Wolverhampton is

⁷ The existence of structure in the residuals of any model will suggest that there is something that is not being taken into account, be it a variable or as in this case a time series behaviour.

increased what happens elsewhere and does this effect last a long time? In this case, as the data is monthly, there is the possibility that a change will have a contemporaneous effect on the other outcomes. This means that an *orthogonal impulse response function* is required (the forecast IRF does not allow us to consider this possibility). The order of the variables has an impact on the IRF so the VARs were estimated 50 times with random ordering of the variables and averaged. Though there are over 40000 (8!) possible combinations, this sample size was considered reasonable to get an overall feel of the impact of the changes.

The next step is to deal with the structure that is in the models previously considered. Given that there is non-stationarity in the series, the (first) differences in the variable of choice are used. This works because the *changes* are normally distributed in a statistically better sense (i.e. the differencing makes them stationary). This is supported using the univariate (ARIMA) where the algorithms and estimation finds that the first difference is important. The autocorrelation within the residuals is removed.




Figure 23 Autocorrelation of VAR is Differences

As before, the changes in violent crime can be forecast up to 5 years hence. The impulse response functions are also presented. This again shows the inter-relationship between areas. Noticeably there is very limited impact in the changes in the levels, though this is not zero across the board the effect diminishes quickly to zero.




Figure 24 Forecast of VAR in Differences (Forecasts are Increases in Crime)

Key: red mean forecast, sky-blue 80th % interval, grey 95th % interval




Figure 25 Five Year Forecast of VAR in Differences

Key: red mean forecast, sky-blue 80th % interval, grey 95th % interval




Figure 26 Impulse Response Functions for VAR in Differences

The model's accuracy for the cross validated VAR is relatively complex to calculate. The in training model's results are presented below. Solihull's odd reported value is due to the prevalence of 0s in the forecasts.

Table 4: Vector Auto-Regression Prediction Fits.

LPU	RMSE	MAE	MAPE
BE	14.012	10.995	21.790
BW	17.408	13.061	17.174
CV	10.458	7.994	20.389
DY	7.896	6.027	23.581
SH	4.999	3.876	Inf
SW	10.564	8.063	25.286
WS	8.108	6.257	30.035
WV	8.964	6.957	24.627
Overall	10.301	7.904	23.269

Table 5: Cross- Validation Accuracies for Levels VAR

	CV Mean RMSE	CV Median RMSE	CV SD RMSE	CV Mean MAE	CV Median MAE	CV SD MAE	CV Mean MAPE	CV Median MAPE	CV SD MAPE
BE	19.015	18.193	8.153	17.182	15.636	8.384	16.223	14.084	7.903
BW	21.763	19.011	10.051	19.016	16.804	9.955	13.757	11.298	7.583
CV	14.020	12.752	7.431	12.687	10.526	7.540	24.533	15.918	20.928
DY	11.677	9.551	6.703	10.848	8.857	6.862	20.850	18.401	11.301
SH	5.991	4.803	4.093	5.525	4.435	4.183	21.328	17.211	13.872
SW	12.919	11.037	8.055	11.643	8.852	8.150	17.652	13.832	10.268
WS	10.200	8.442	6.211	9.259	6.789	6.320	17.617	14.486	10.223
WV	12.153	9.930	7.457	11.191	8.556	7.451	18.653	17.600	9.228
Overall	13.467	11.481	8.729	12.169	10.081	8.476	18.827	14.611	12.429

The cross validation was also performed on the VAR in differences. The results are given below. There was a degree of variation across the model space, but the majority of the cross validated models (80%) are VAR (3) models implying that an order 3 lag is employed when modelling the levels. The log-likelihoods etc. are uniformly distributed approximated with relatively little spread in the values. The residual variances are given below and show the importance of the multivariate models with the non-diagonal elements *not* being equal to zero.

Table 6: Cross- Validation Residual Variances and Standard Deviations for Levels VAR

	BE	BW	CV	DY	SH	SW	WS	WV
BE	208.194							
BW	95.355	335.281						
CV	43.565	66.311	110.552					
DY	26.070	27.994	9.106	54.727				
SH	14.978	11.876	14.412	1.008	24.311			
SW	46.762	47.598	28.699	19.475	13.313	109.021		
WS	34.711	38.115	13.365	6.457	9.409	8.050	69.510	
WV	32.434	48.855	21.528	14.430	6.365	7.853	12.563	79.963
SD(BE)	4.337							
SD(BW)	2.991	7.634						
SD(CV)	1.803	5.634	8.896					
SD(DY)	4.378	4.196	2.572	5.7685				
SD(SH)	2.870	2.166	1.138	0.833	1.520			
SD(SW)	1.757	8.089	3.468	3.895	1.701	6.582		
SD(WS)	2.253	4.757	3.194	3.154	0.637	1.588	3.855	
SD(WV)	3.947	4.301	2.875	2.871	0.867	1.374	3.863	5.949

7.2.2 Auto-Regressive Distributed Lag Model

One of the challenges associated with the VAR models is that they are simply based on the series in question and are seen by some as atheoretical. In the situation where we can have variables informed by the literature, it could make sense to use this (at least to see if predictions are improved). Further the ARDL approach allows us to mix data integrated with different orders (qv Pesaran, Shin, and others (1995)). This means that we are more directly able to model the variables without an extra level of testing for model specification. The residuals are considered for these models to ensure that there is no remaining evidence of a unit-root⁸ or serial correlation. The models were also run in error correction form, where the dependent variable is run as a difference however this makes no significant difference. The model was estimated using a data set without missing values. The data following that was used as the test dataset as it allowed us to test the stacking process as well. It did however limit the window for the forecast to a year.

⁸ A unit root would mean non-stationarity (not necessarily due to a trend).

Table 7: Covariance matrix and Correlation Matrix for Auto-Regressive Distributed Lag Models Residuals

	BE	BW	CV	DY	SH	SW	WS	WV
BE	103.381							
BW	32.920	131.523						
CV	15.512	5.498	37.202					
DY	21.527	6.294	3.412	33.584				
SH	17.435	7.09	9.946	0.68	17.459			
SW	16.420	16.762	14.758	-3.3	10.417	47.568		
WS	22.402	22.428	8.475	5.893	5.645	9.716	44.778	
WV	3.827	5.104	0.815	6.741	-0.647	-3.504	4.725	40.346
BE	1.000							
BW	0.282	1						
CV	0.250	0.079	1					
DY	0.365	0.095	0.097	1				
SH	0.410	0.148	0.39	0.028	1			
SW	0.234	0.212	0.351	-0.083	0.361	1		
WS	0.329	0.292	0.208	0.152	0.202	0.211	1	
WV	0.059	0.07	0.021	0.183	-0.024	-0.08	0.111	1

Table 8 ARDL Model Specification Metrics

Model	Log_lik	AIC	AICc	BIC
Overall	-270.706	627.413	729.683	730.374

As before, there are significant areas where there is a correlation between the residuals suggesting a geo-spatial element that is to some extent picked up in the vector specification. In two cases using the Breusch Godfrey test for auto-correlation there is no evidence of first order auto-correlation, though there may be some of higher order. Note that the likelihoods and information criteria are not comparable to the VAR as the sample is considerably different. These are provided for completeness.

We can use simulations to consider the impact, given the simplest forms of the models here to, of a number of variables. An example of an increase in the precept in Birmingham (standardized), an increase in the proportion of council tax outstanding (in Coventry) and a third on the impact of increases in non-working households (in Wolverhampton) is used to show the impact on violent crimes. The graphs show the impact on the relevant areas. The precept changes in Birmingham have different impacts across the city. Increases in the precept (as standardized) see a reduction in violence in the West but not in the East. The simulations involved a shock to the

relevant variable 10 months into the series. This was a positive shock of 1 standard deviation. It gives the same sort of information as the impulse response function but in a different manner.

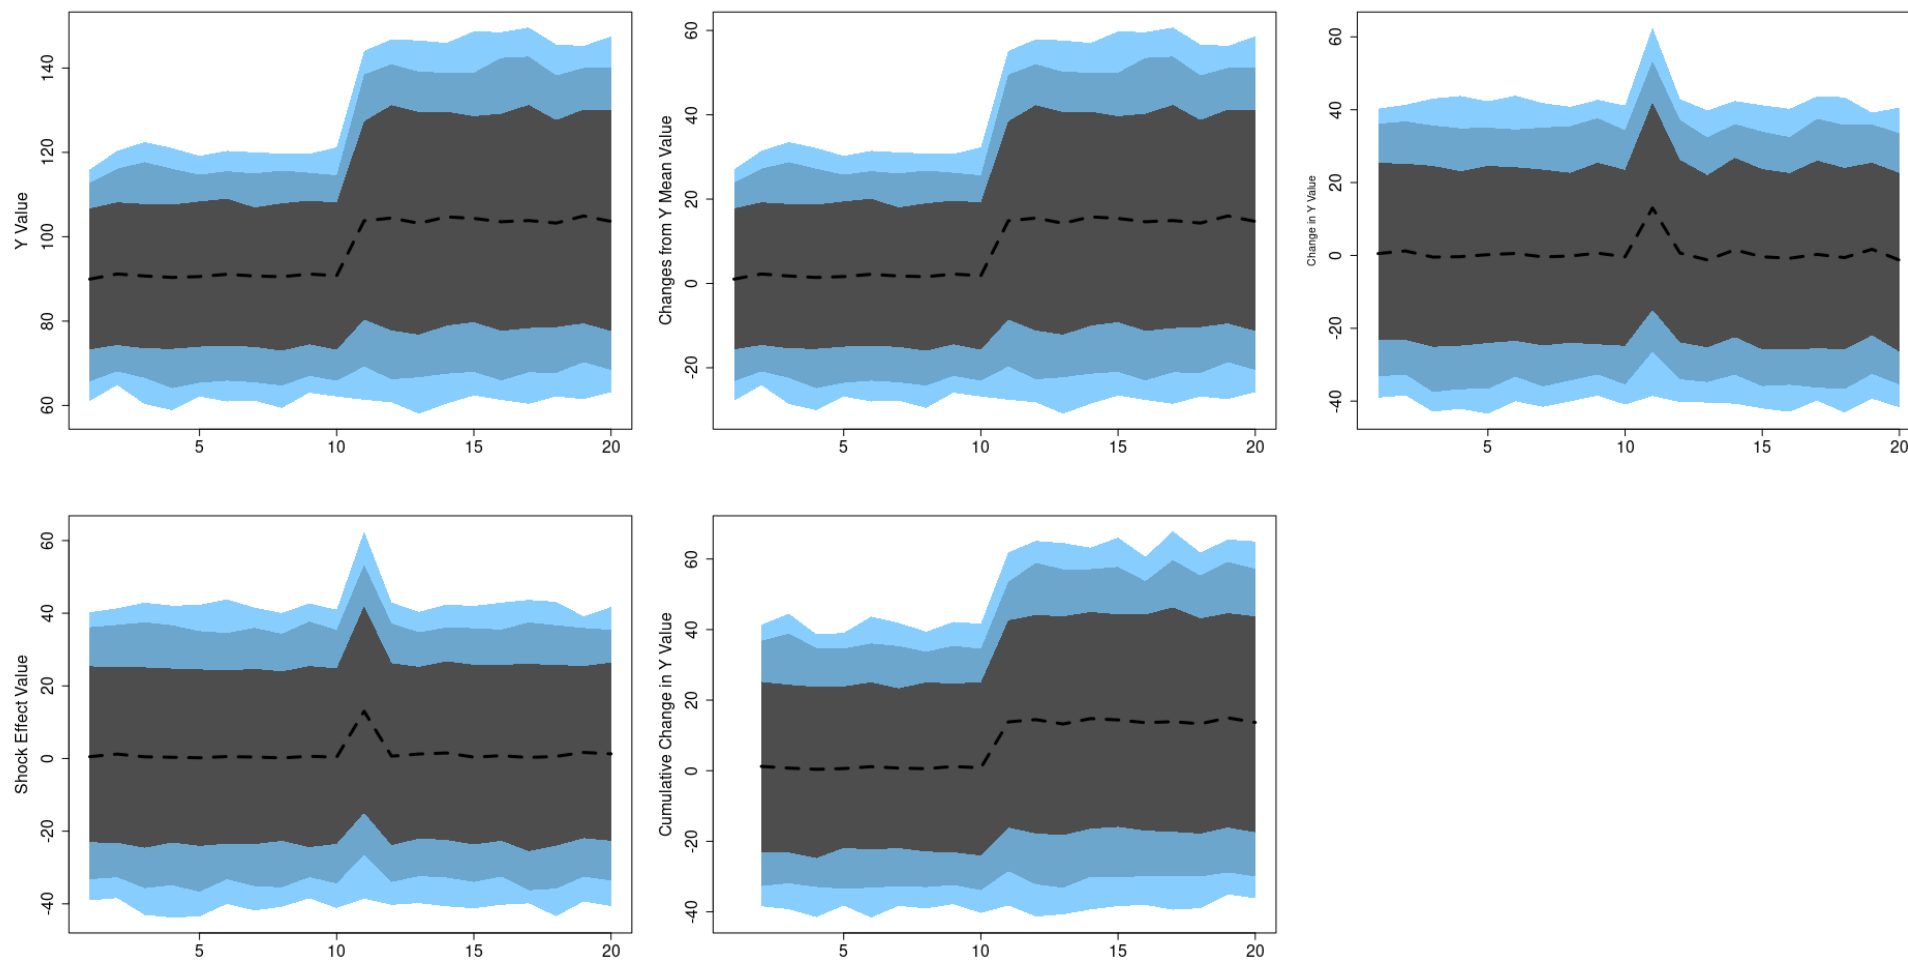


Figure 27 Simulation Effects of Changes to Precept in Birmingham on Birmingham East

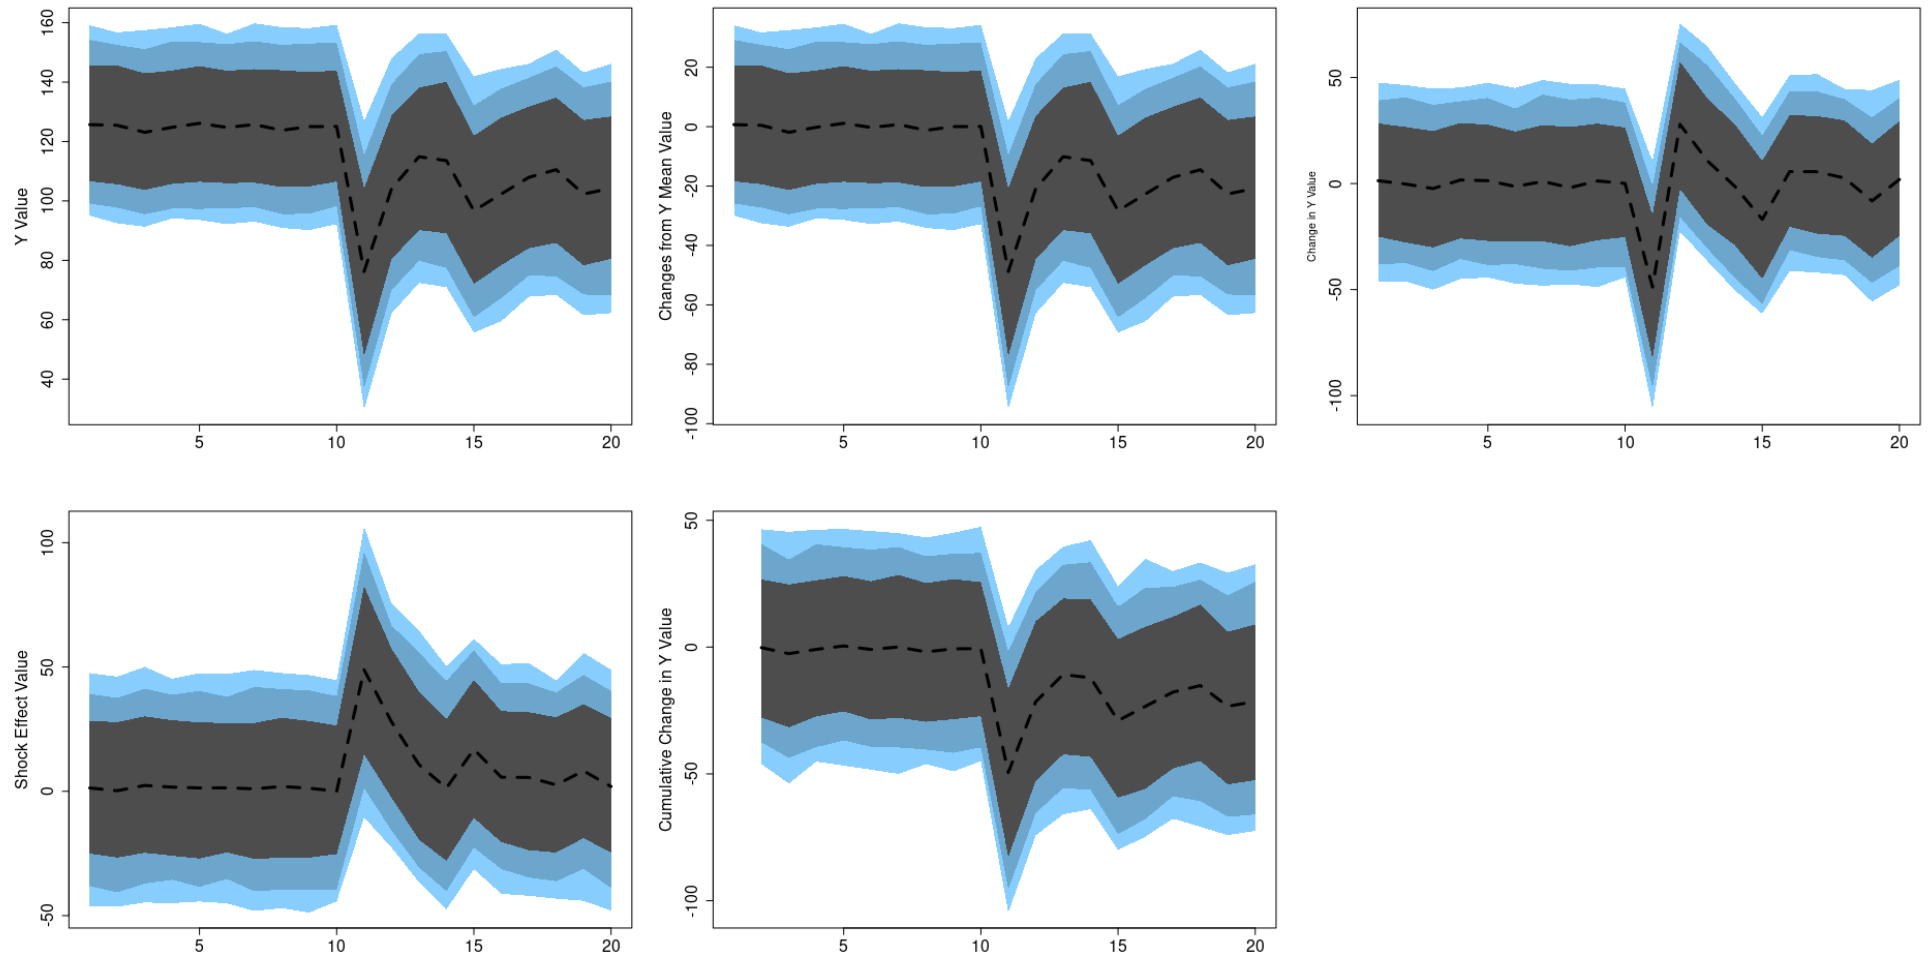


Figure 28 Simulation Effects of Changes to Precept in Birmingham on Birmingham West

An increase in the council tax nonpayment in Coventry has a positive impact on violence. This is *not* saying that there is causality in any way, rather that there is a relationship where the increase in inability to pay the tax is associated with an environment that has higher levels of violence. It is reflective of the social situation where people find themselves unable to pay their council tax.

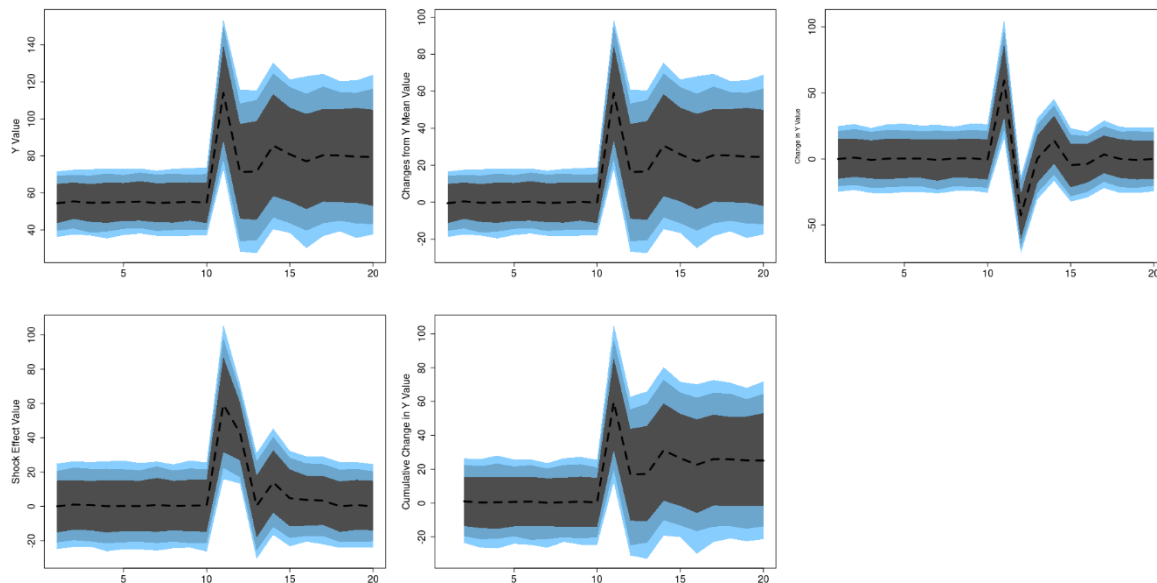


Figure 29 Simulation Effects of Changes to Non-Payment of Council Tax in Coventry on Coventry

Looking at the (un)employment levels in Wolverhampton, we see that there is little difference. There is a very short term increase but this then returns to the previous levels of violence in Wolverhampton; there is a dip in violence in Birmingham East and Birmingham West. The cause of this is not clear but the fact that there is an effect across Birmingham suggests that the increase in non-working households in Wolverhampton might be reflecting some substitution with workers in Birmingham.

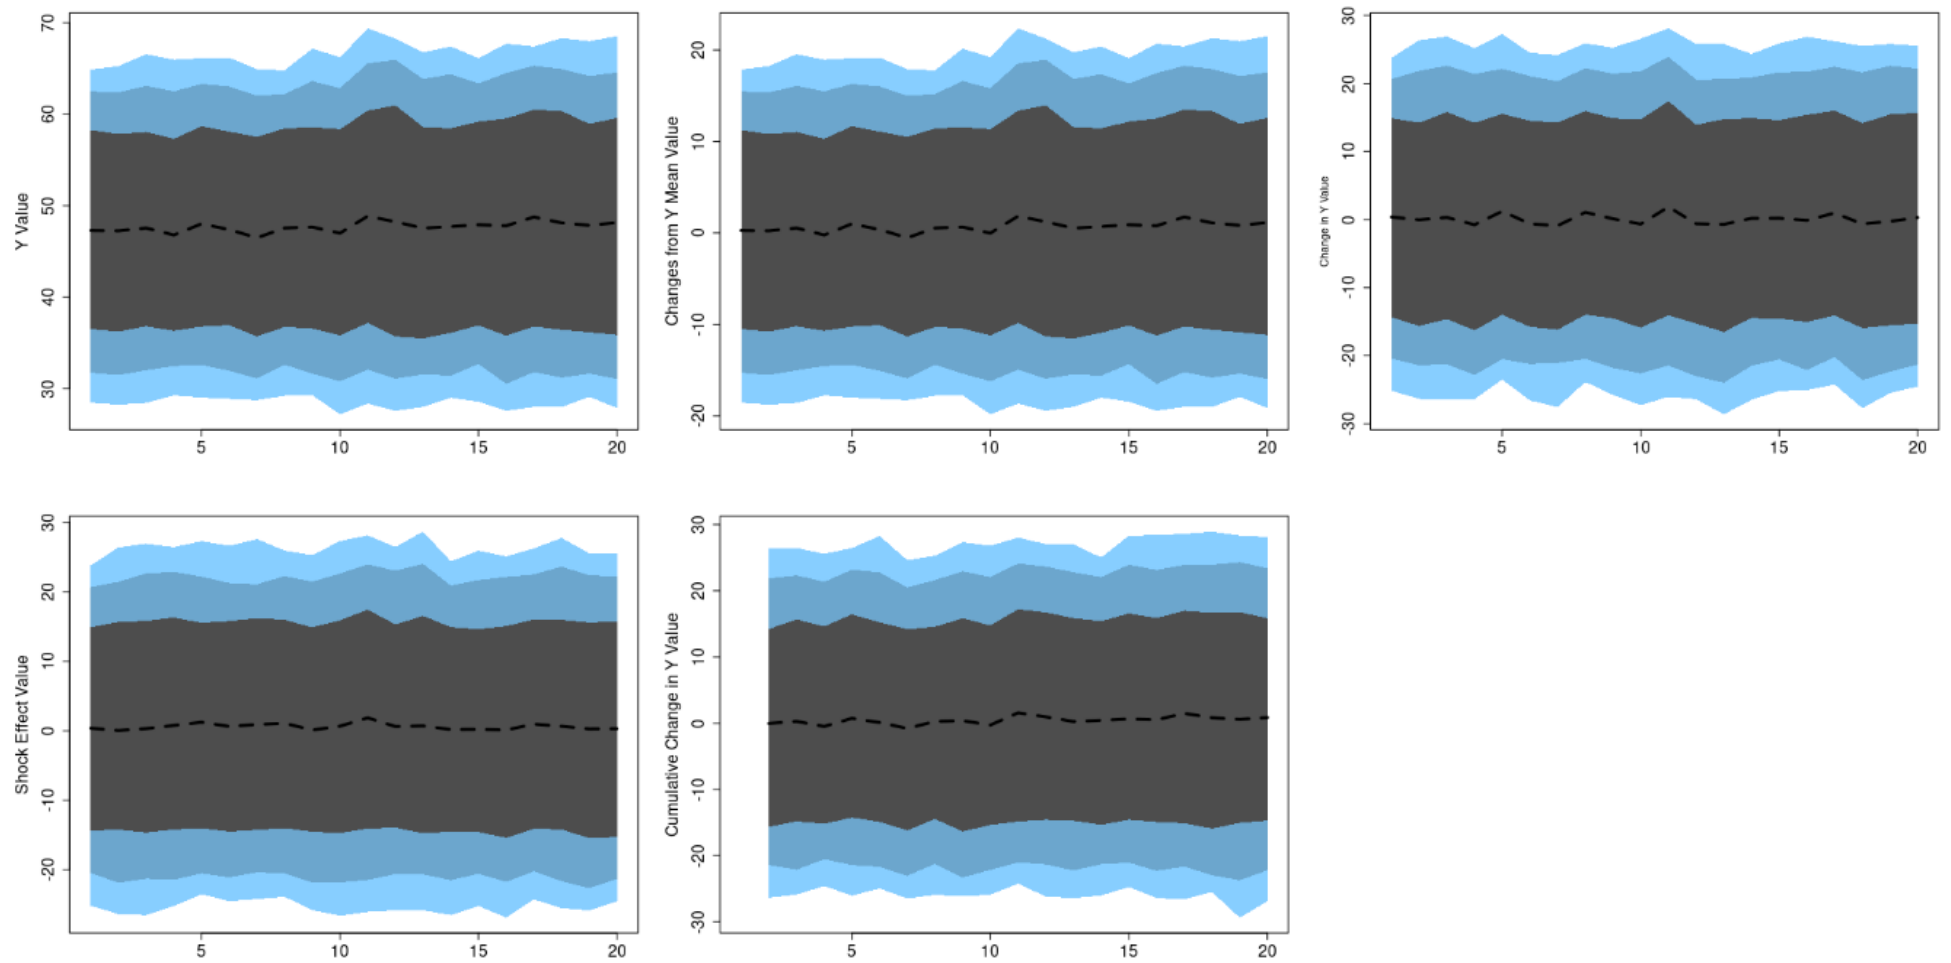


Figure 30 Simulation Effects of Changes in Non-working households in Wolverhampton on Violence in Wolverhampton

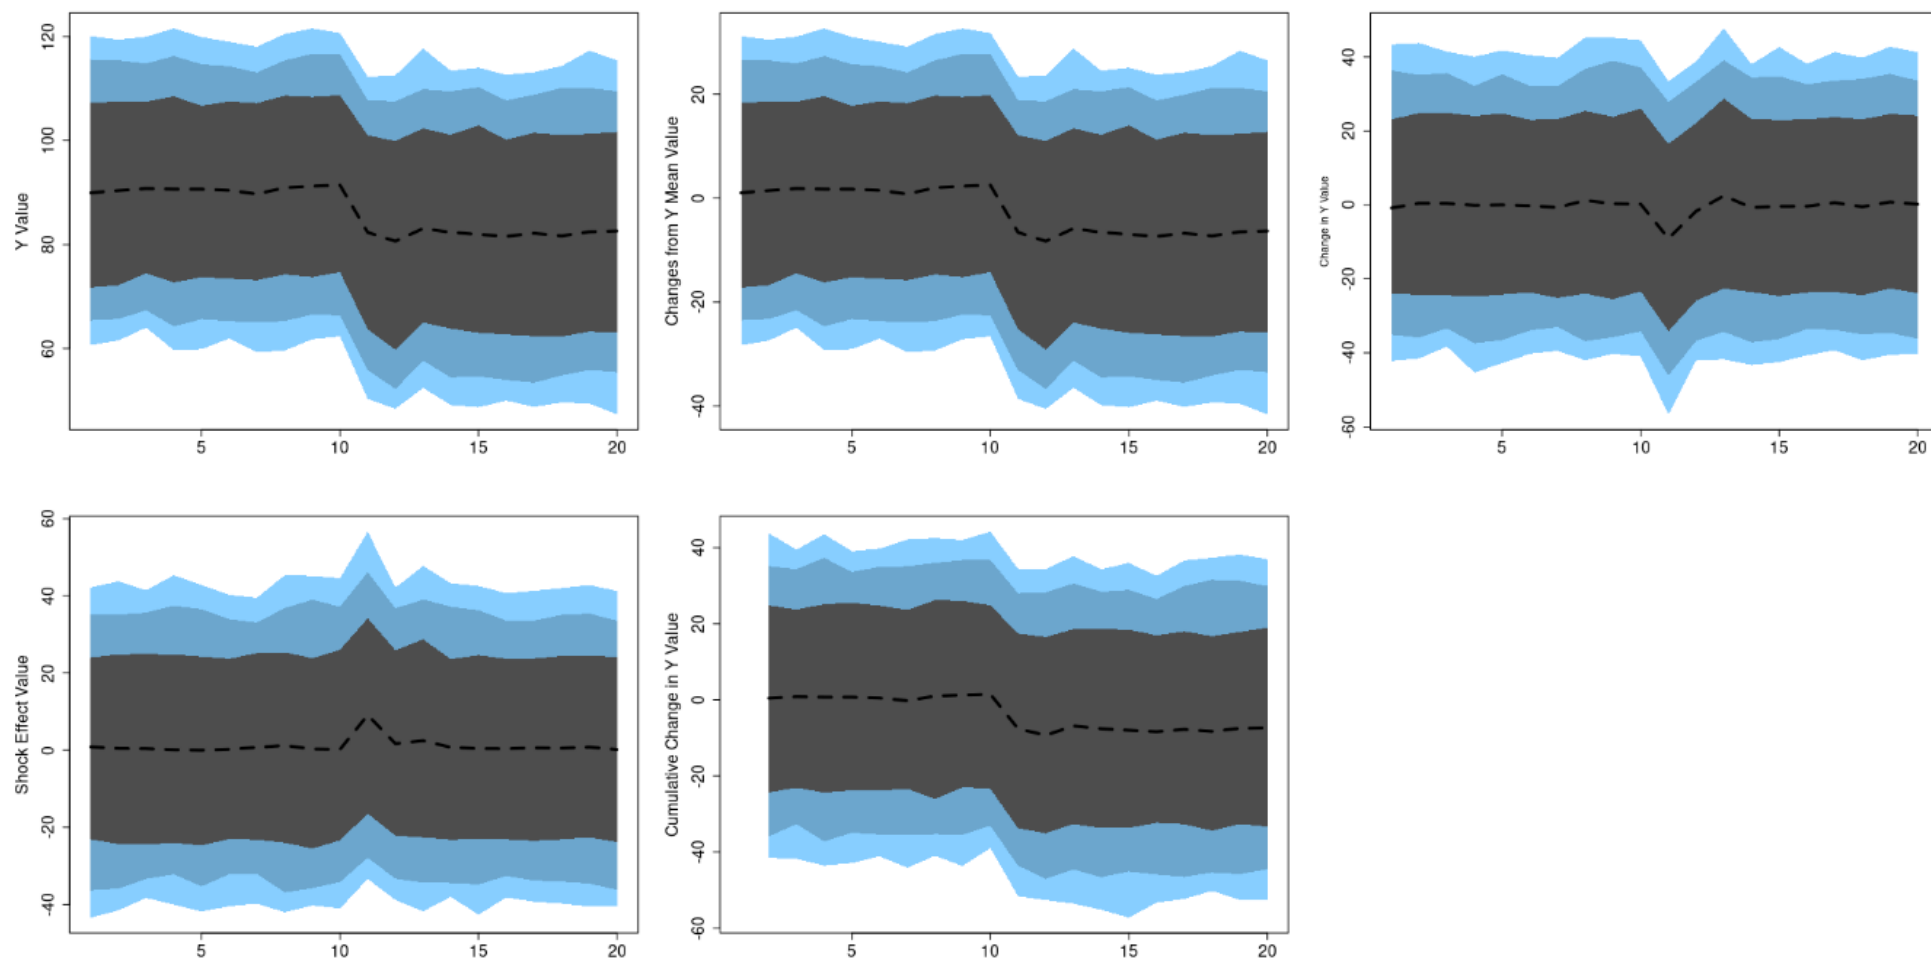


Figure 31 Simulation Effects of Changes in Non-working households in Wolverhampton on Violence in Birmingham East

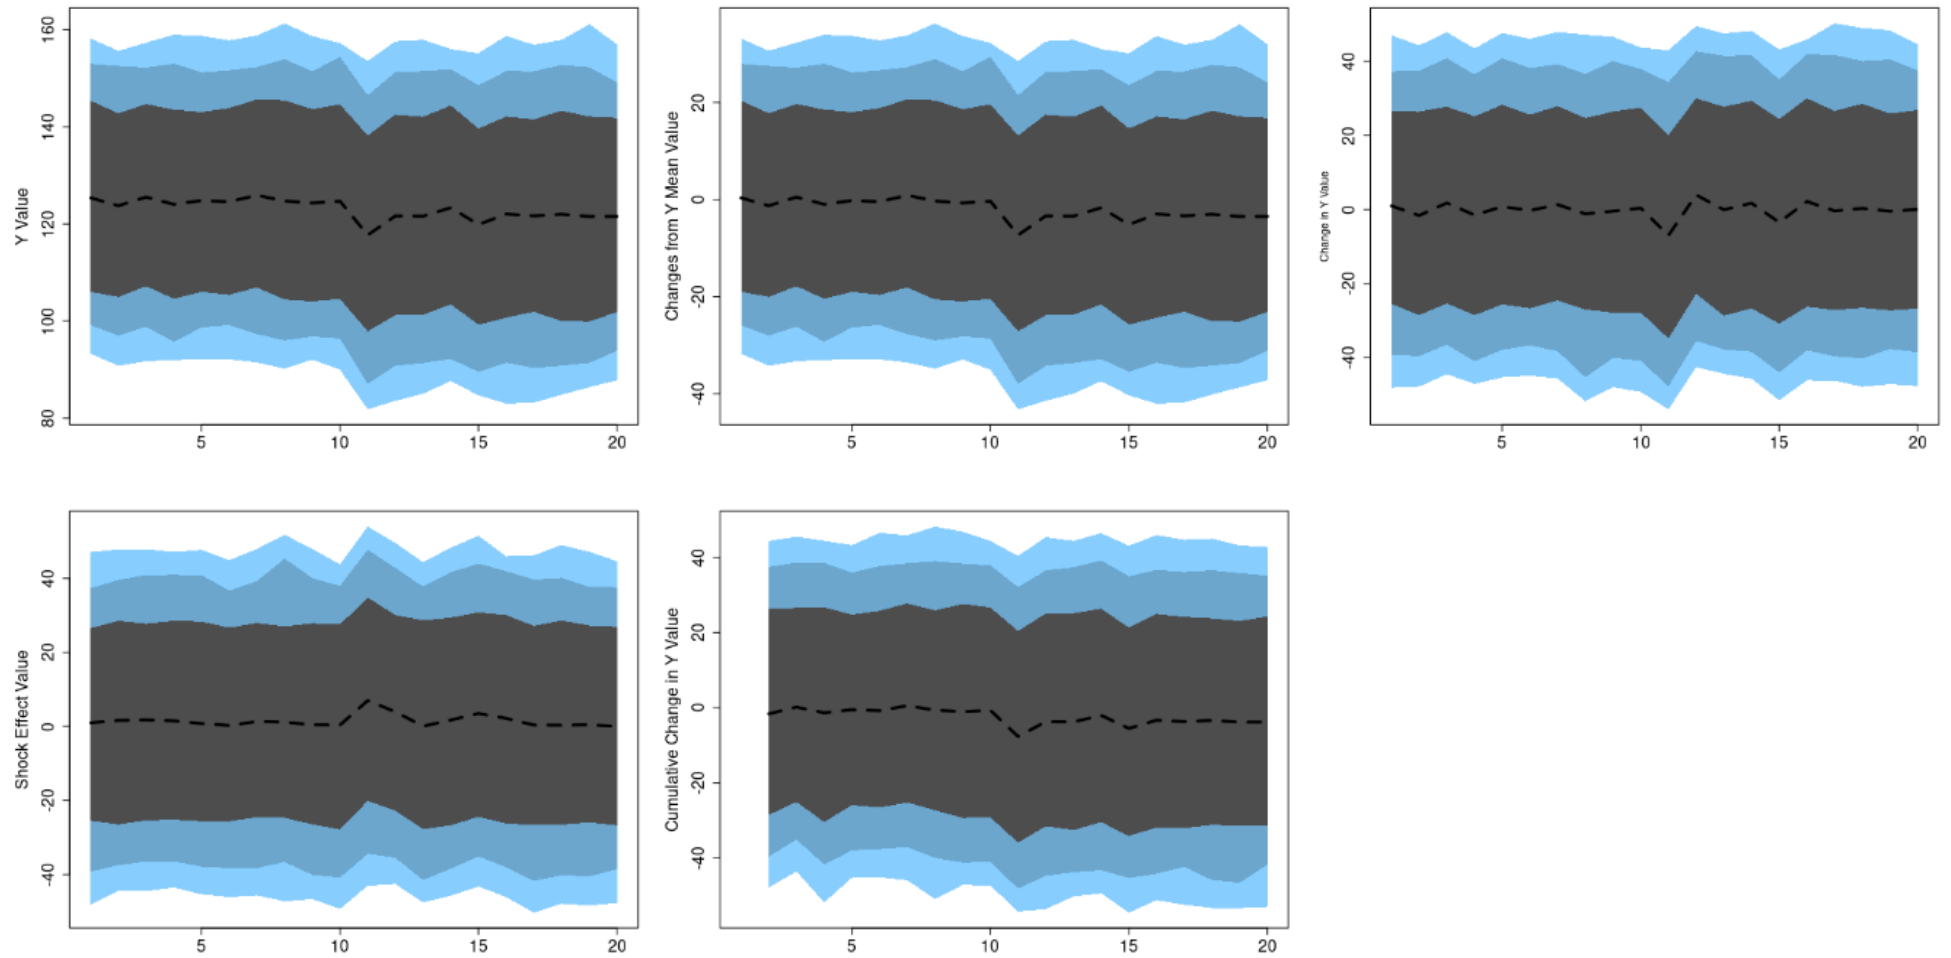


Figure 32 Simulation Effects of Changes in Non-working households in Wolverhampton on Violence in Birmingham West

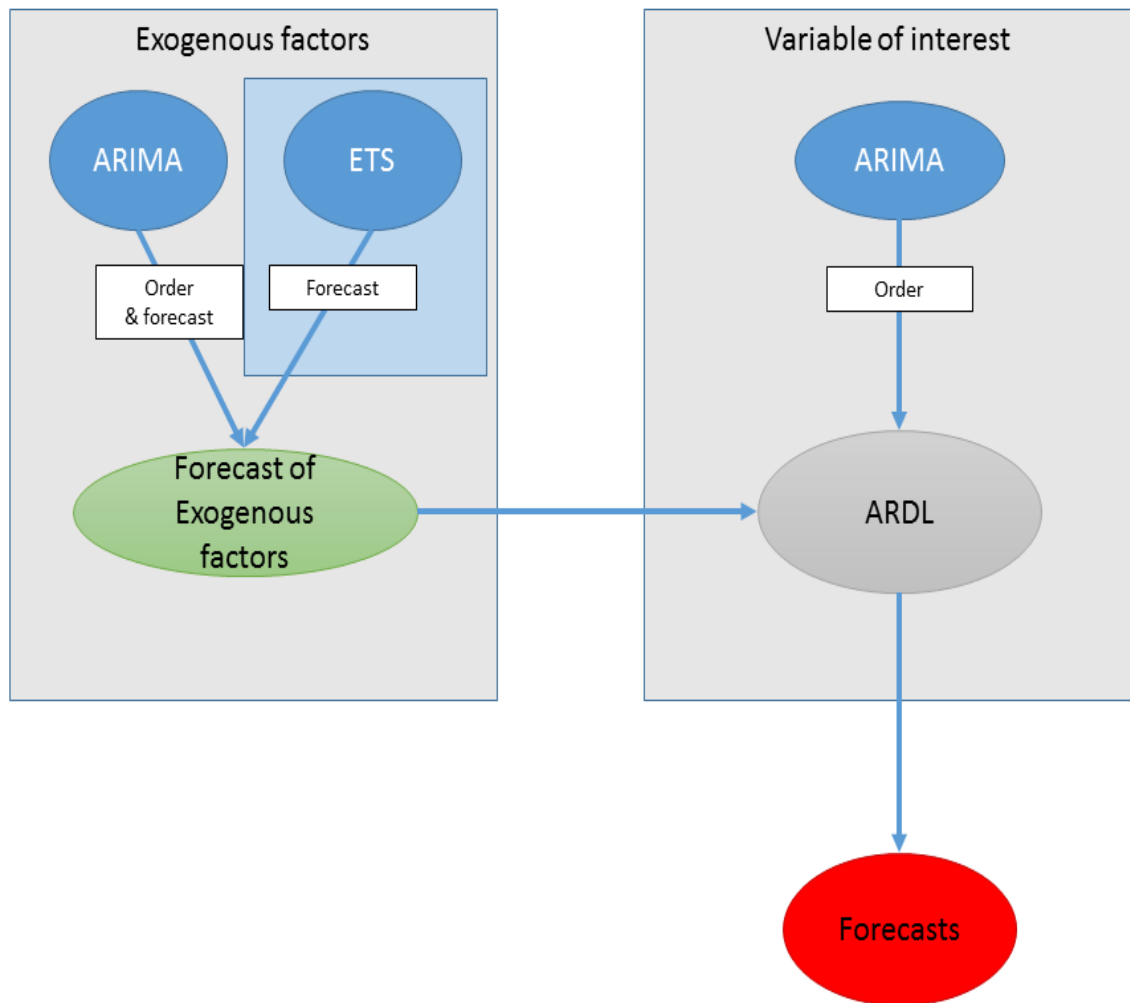


Figure 33 Approach to Stacked Forecasting Within ARDL Framework

Forecasting off these models will be complicated. Not all of the data is available for all the time period of interest. It is therefore necessary to use a combination of the forecasts and the actual outcomes in the modelling. The model is forecast off the complete data; with a pseudo-test set to test the forecasts of the missing data within the model's framework. The main variables without a full term of data are the truancy levels (ends 2019-08-01), free meals for secondary school students (2020-08-01 based on school year) and non- working households (end of 2019). The more financial factors such as council tax non-payment all run until April 2020. This is due to the financial year reporting regulations. The restricted data begins in 2012-09-01 and runs to 2019-08-01. In the first case the truancy forecasts are added to the dataset as these stop at that point. From this point on the forecasts of the absences from schools are included. From January 2020, the non-working data is missing. This is then replaced by the forecasts. In May 2020, the precept and the council tax failure to pay are replaced by forecasts. Finally in September 2020, the free school meals are replaced by their forecasts.

The final data available ends in March 2021, the data can be squared with the forecasts filling in the gaps starting in 2012. Though not a long series, using the ARDL estimated until August 2019 this gives about 18 months to forecast where we know the outcomes and are able to sensibly judge the model. In the first case, ARIMA models were used to fill in the gaps (so as to keep the modelling techniques similar. If the ARIMA approach is not successful then ETS data can be introduced). The model estimated on the pre-filled data will be used to roll forwards with the data that contains the missing values. It will only use data that we would know at that time.

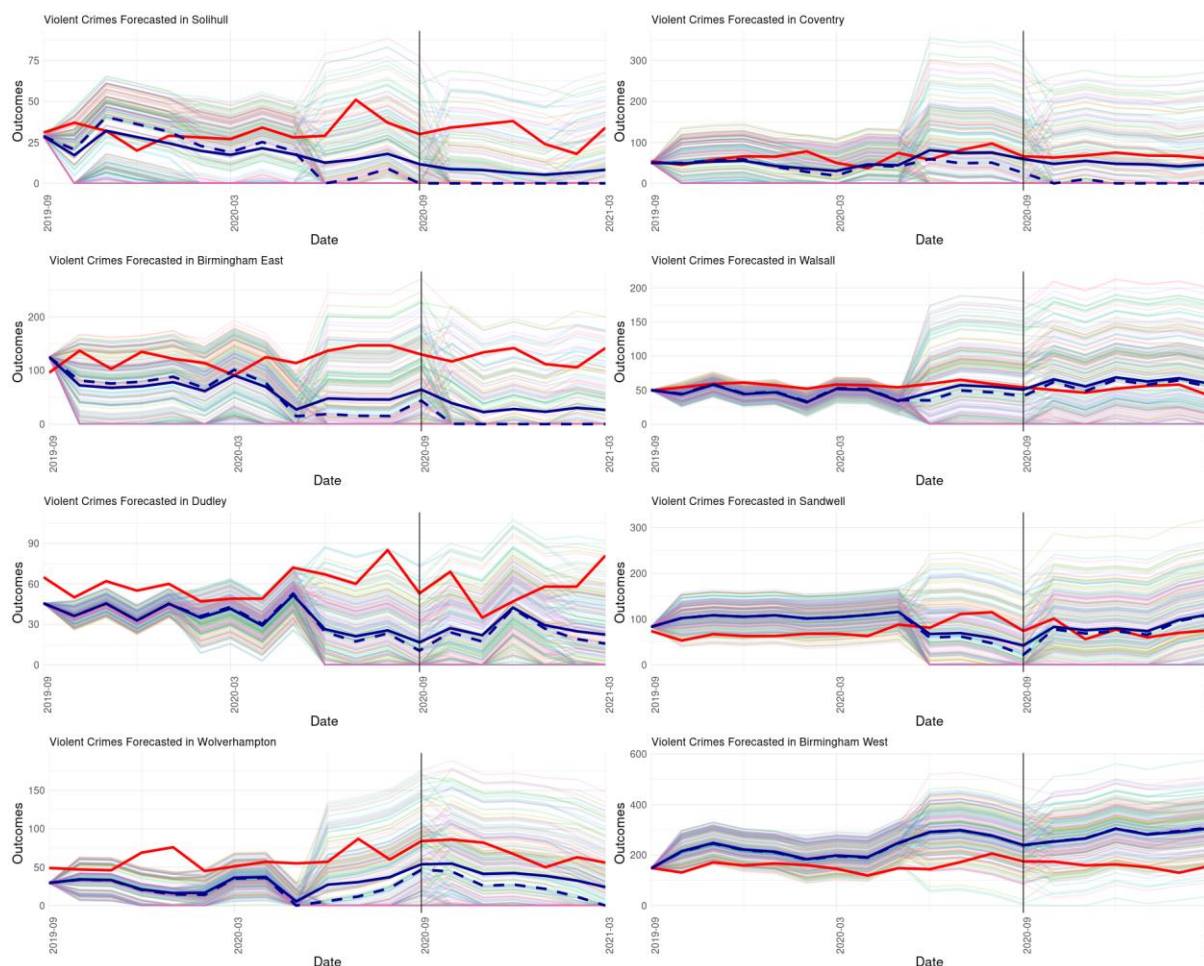


Figure 34 Forecasts of Violence from ARDL model

The averaged forecasts (based on 200 simulations) in solid blue (the median is dashed blue) in the figure above show a lot of variation; compared to the red line of the actual outcomes. The forecasts are quite well clustered until 9 months ahead, when the spread increases considerably. There is no impact of Covid explicitly in the models. There was not enough data in the modelling (or training) data to include this in the ARDL model to ascertain the impact of the pandemic. This would obviously be a problem with the model but, despite that the mean is generally reflective of the actual outcomes.

In examining the outputs and accuracy of the models, **it is clear that the impact of the independent variables is not overly beneficial.** In the model building phase, actual data was used to estimate the model. This involved using forecasts of these variables, based on relatively short annual data expanded into monthly data. The forecasts generated off these extra forecasts is adding significant extra uncertainty. The short

data series used in the forecasted values of the independent variables are most likely not well specified and thus the projections off these are at best hopeful and most likely of variable use. The use of multiple paths from these forecasts does allow some understanding of these uncertainties. The uncertainty associated with them suggests that the better specification and greater informational certainty of the Vector Autoregression is a more sound footing.

Table 9: Vector Auto-Regressive Distributed Lag Models.

	RMSE	MAE	MAPE
BE	78.445	72.263	56.723
BW	99.379	89.363	58.392
CV	19.782	17.029	25.757
DY	30.651	26.377	42.557
SH	18.857	16.115	49.894
SW	34.290	31.059	43.374
WS	11.879	10.324	19.281
WV	33.440	30.240	47.433
Overall	40.840	36.596	42.927

As expected the errors from the forecasts are quite large relative to those of the VAR. This is due to the increased uncertainty associated with the models due to the stacking of the forecasts and the frequency matching in the modelling stages. There is a general directional correspondence and there is explanatory power in the ARDL model, however it is not good at forecasting over the year.

There are some areas of success; Coventry, Walsall and Solihull and to a lesser extent Wolverhampton look reasonable forecasts, though Wolverhampton looks as if there is a systematic under-estimation. These are beacons that are too easy to over-emphasize. **The models degrade once the additional data is stacked into the models, suggesting that the increased predictive error is potentially too costly in a modeling sense to be useful.**

8 Conclusions & Recommendations

Though it was a sensible route to investigate, the outcome of including extra explanatory variables added little value with the data in the current state of specifically the local authority information being annual and the short-time scale of the data. The data might be useful at a later date and holds some potential value as demonstrated when the data was more certain. **They currently have little predictive power and might even reduce the power of the simply autoregressive processes suggested elsewhere.**

The simpler models have some predictive power though as would be expected the accuracy of these models reduces as the time horizon increases. The predictive intervals are broad and associated with the relatively high levels of variability in the data and for example the regime shifts caused by COVID-19.

The additional explanatory data might be helpful in some aspects and areas however currently the data available is coarse and short which limits the usefulness in the prediction of the violent crimes. Some of these variables are themselves difficult to forecast and prone to error and uncertainty and this adds to the final levels of predictive uncertainty associated with the forecasts.

9 Appendix

9.1.1 VAR Coefficients

Coefficients for VAR in levels

	Estimate BE	Std.error BE	Estimate BW	Std.error BW	Estimate CV	Std.error CV	Estimate DY	Std.error DY	Estimate SH	Std.error SH	Estimate SW	Std.error SW	Estimate WS	Std.error WS	Estimate WV	Std.error WV
Lag 1 BE	0.130	0.075	0.257	0.093	0.138	0.056	0.034	0.042	0.050	0.027	0.035	0.056	0.028	0.043	0.104	0.048
Lag 2 BE	0.113	0.075	0.055	0.094	0.011	0.056	0.006	0.042	-0.022	0.027	0.064	0.057	0.030	0.044	0.053	0.048
Lag 3 BE	0.030	0.075	-0.044	0.093	-0.061	0.056	0.018	0.042	0.000	0.027	0.042	0.057	0.039	0.043	0.011	0.048
Lag 1 BW	0.251	0.061	0.377	0.076	0.174	0.046	0.058	0.034	0.056	0.022	0.075	0.046	0.056	0.035	0.031	0.039
Lag 2 BW	0.047	0.066	0.091	0.083	0.002	0.050	0.019	0.037	-0.001	0.024	-0.006	0.050	0.048	0.038	0.069	0.043
Lag 3 BW	0.086	0.062	0.045	0.077	0.066	0.046	-0.071	0.035	0.019	0.022	-0.032	0.046	-0.046	0.036	-0.070	0.039
Lag 1 CV	-0.067	0.094	0.251	0.117	0.096	0.070	-0.074	0.053	-0.036	0.034	0.003	0.071	-0.085	0.055	-0.001	0.060
Lag 2 CV	-0.181	0.095	-0.060	0.118	0.159	0.071	-0.087	0.053	0.017	0.034	0.077	0.071	-0.031	0.055	-0.045	0.061
Lag 3 CV	-0.056	0.096	0.126	0.119	0.222	0.071	0.111	0.054	0.004	0.034	-0.001	0.072	0.071	0.055	-0.009	0.061
Lag 1 DY	0.069	0.129	-0.148	0.160	0.117	0.096	0.156	0.072	-0.059	0.046	0.079	0.097	-0.034	0.074	-0.020	0.082
Lag 2 DY	0.364	0.126	0.260	0.157	0.057	0.094	0.199	0.071	0.063	0.045	0.221	0.095	0.212	0.073	0.081	0.081
Lag 3 DY	-0.127	0.128	0.182	0.158	-0.200	0.095	-0.072	0.072	-0.010	0.045	-0.137	0.096	-0.018	0.074	0.041	0.082
Lag 1 SH	0.306	0.203	0.514	0.252	0.070	0.152	0.058	0.114	0.280	0.072	0.080	0.153	0.221	0.117	-0.119	0.130
Lag 2 SH	0.025	0.212	-0.211	0.264	-0.340	0.158	-0.203	0.120	-0.079	0.076	-0.250	0.160	0.065	0.123	-0.016	0.136
Lag 3 SH	0.222	0.202	0.089	0.251	0.156	0.151	0.146	0.114	0.124	0.072	0.279	0.152	-0.079	0.117	0.085	0.129
Lag 1 SW	0.093	0.099	0.146	0.123	0.066	0.074	0.127	0.056	0.001	0.035	0.220	0.075	-0.007	0.057	0.206	0.063
Lag 2 SW	0.023	0.102	0.096	0.127	0.026	0.076	0.043	0.058	0.070	0.036	0.049	0.077	-0.065	0.059	0.145	0.065
Lag 3 SW	-0.144	0.104	-0.187	0.129	0.016	0.077	-0.027	0.058	-0.018	0.037	0.091	0.078	0.021	0.060	-0.101	0.066
Lag 1 WS	0.017	0.121	-0.276	0.150	-0.080	0.090	0.102	0.068	0.075	0.043	0.168	0.091	0.286	0.070	-0.014	0.077
Lag 2 WS	0.057	0.127	0.089	0.158	0.078	0.095	0.142	0.071	-0.010	0.045	-0.020	0.096	0.004	0.073	-0.060	0.081
Lag 3 WS	-0.013	0.123	0.185	0.153	-0.111	0.092	0.096	0.070	-0.024	0.044	-0.008	0.093	0.142	0.071	0.033	0.079

Lag 1 WV	0.137	0.108	-0.020	0.134	0.041	0.081	0.049	0.061	0.036	0.039	-0.004	0.082	0.005	0.063	0.166	0.069
Lag 2 WV	-0.093	0.108	0.007	0.134	-0.130	0.080	0.024	0.061	-0.045	0.038	-0.072	0.081	0.002	0.062	0.029	0.069
Lag 3 WV	0.128	0.104	-0.077	0.129	-0.042	0.078	0.022	0.059	0.009	0.037	0.051	0.078	0.031	0.060	0.155	0.066

Coefficients for VAR in Differences

	Estimate BE	Std.error BE	Estimate BW	Std.error BW	Estimate CV	Std.error CV	Estimate DY	Std.error DY	Estimate SH	Std.error SH	Estimate SW	Std.error SW	Estimate WS	Std.error WS	Estimate WV	Std.error WV
Lag 1 BE	-0.667	0.071	0.205	0.084	0.165	0.051	0.014	0.040	0.039	0.026	0.033	0.05183768	0.024	0.039	0.071	0.046
Lag 2 BE	-0.343	0.084	0.267	0.099	0.206	0.060	-0.004	0.047	-0.001	0.030	0.082	0.06100877	0.032	0.046	0.087	0.054
Lag 3 BE	-0.121	0.071	0.174	0.084	0.131	0.051	-0.003	0.040	0.002	0.026	0.105	0.05190693	0.040	0.039	0.044	0.046
Lag 1 BW	0.174	0.061	-0.508	0.072	0.113	0.044	0.028	0.035	0.037	0.022	0.015	0.04445529	0.005	0.034	0.027	0.039
Lag 2 BW	0.100	0.066	-0.296	0.078	0.064	0.048	0.016	0.038	0.010	0.024	-0.038	0.04835856	0.024	0.037	0.097	0.043
Lag 3 BW	0.067	0.060	-0.204	0.071	0.046	0.043	-0.063	0.034	0.016	0.022	-0.062	0.04346772	-0.029	0.033	0.020	0.039
Lag 1 CV	-0.048	0.098	0.100	0.116	-0.841	0.071	-0.016	0.056	-0.009	0.035	0.060	0.07158	-0.011	0.054	0.035	0.064
Lag 2 CV	-0.167	0.116	-0.057	0.137	-0.587	0.083	-0.035	0.065	0.033	0.042	0.198	0.08437458	0.010	0.064	0.011	0.075
Lag 3 CV	-0.124	0.098	-0.046	0.116	-0.264	0.071	0.118	0.055	0.054	0.035	0.221	0.07138793	0.089	0.054	-0.004	0.063
Lag 1 DY	0.028	0.122	-0.218	0.144	0.094	0.088	-0.581	0.069	-0.055	0.044	0.082	0.0887136	-0.087	0.067	-0.064	0.079
Lag 2 DY	0.354	0.139	-0.006	0.165	0.184	0.100	-0.260	0.079	0.015	0.050	0.152	0.10158574	-0.008	0.077	-0.024	0.090
Lag 3 DY	0.141	0.124	0.126	0.147	0.010	0.090	-0.207	0.070	0.009	0.045	-0.127	0.09076037	-0.121	0.069	-0.016	0.081
Lag 1 SH	0.306	0.194	0.335	0.229	0.086	0.140	0.064	0.110	-0.528	0.070	-0.020	0.14135507	0.183	0.108	-0.142	0.125
Lag 2 SH	0.196	0.207	-0.010	0.245	-0.209	0.149	-0.117	0.117	-0.463	0.075	-0.361	0.15083985	0.237	0.115	-0.119	0.134
Lag 3 SH	0.425	0.195	-0.207	0.231	-0.001	0.140	0.081	0.110	-0.156	0.070	0.009	0.14217255	0.099	0.108	-0.053	0.126
Lag 1 SW	0.042	0.097	0.120	0.115	0.069	0.070	0.085	0.055	-0.022	0.035	-0.611	0.07083756	0.004	0.054	0.177	0.063
Lag 2 SW	0.043	0.106	0.174	0.126	0.070	0.077	0.073	0.060	0.046	0.038	-0.406	0.07747537	-0.055	0.059	0.261	0.069
Lag 3 SW	-0.090	0.100	-0.051	0.119	0.061	0.072	-0.001	0.057	-0.001	0.036	-0.166	0.07315132	-0.022	0.056	0.063	0.065
Lag 1 WS	0.037	0.120	-0.263	0.142	-0.051	0.086	0.051	0.068	0.058	0.043	0.180	0.08726822	-0.590	0.066	-0.041	0.077
Lag 2 WS	0.150	0.127	-0.192	0.150	0.039	0.092	0.104	0.072	0.029	0.046	0.099	0.09271437	-0.512	0.071	-0.119	0.082
Lag 3 WS	0.232	0.120	0.056	0.142	-0.040	0.087	0.107	0.068	-0.028	0.043	0.018	0.08762773	-0.308	0.067	-0.069	0.078
Lag 1 WV	0.101	0.107	0.000	0.127	0.066	0.077	0.015	0.061	0.052	0.039	0.020	0.07829484	0.001	0.060	-0.670	0.069
Lag 2 WV	-0.052	0.116	-0.030	0.137	-0.037	0.083	-0.010	0.065	0.015	0.042	-0.049	0.08436775	-0.014	0.064	-0.480	0.075

Lag 3 WV	0.023	0.104	-0.162	0.123	-0.053	0.075	-0.042	0.059	0.043	0.038	0.082	0.076	-0.003	0.058	-0.198	0.067
----------	-------	-------	--------	-------	--------	-------	--------	-------	-------	-------	-------	-------	--------	-------	--------	-------

9.1.2 Auto Regressive Distributed Lag Coefficients

	Estimate BE	Std. Error BE	Estimate BW	Std. Error BW	Estimate CV	Std. Error CV	Estimate DY	Std. Error DY	Estimate SH	Std. Error SH	Estimate SW	Std. Error SW	Estimate WS	Std. Error WS	Estimate WV	Std. Error WV
(Intercept)	-3119.8	1568.21	5773.913	1768.829	1188.507	940.740	866.636	893.822	-668.27	644.451	2488.431	1063.759	719.452	1032.09	-387.368	979.682
Lag 1.BE	0.026	0.174	0.118	0.196	-0.061	0.104	0.178	0.099	0.106	0.072	-0.052	0.118	0.114	0.115	0.127	0.109
Lag 2.BE	-0.126	0.153	0.219	0.172	0.044	0.092	0.023	0.087	0.036	0.063	0.097	0.104	0.071	0.101	0.052	0.095
Lag 3.BE	0.057	0.160	0.385	0.181	-0.029	0.096	0.042	0.091	0.069	0.066	0.148	0.109	0.099	0.105	0.029	0.100
Lag 1.BW	0.273	0.177	-0.559	0.200	-0.203	0.106	-0.054	0.101	-0.003	0.073	-0.168	0.120	-0.050	0.117	0.093	0.111
Lag 2.BW	0.127	0.159	-0.550	0.179	-0.094	0.095	-0.042	0.091	-0.092	0.065	-0.195	0.108	0.075	0.105	0.133	0.099
Lag 3.BW	0.048	0.152	-0.424	0.172	-0.027	0.091	-0.105	0.087	0.052	0.063	-0.077	0.103	-0.071	0.100	0.092	0.095
Lag 1.CV	0.423	0.252	-0.239	0.284	-0.731	0.151	-0.071	0.144	-0.177	0.104	-0.048	0.171	-0.148	0.166	-0.041	0.157
Lag 2.CV	-0.241	0.276	-0.552	0.311	-0.514	0.166	-0.354	0.157	-0.074	0.113	0.044	0.187	0.035	0.182	-0.013	0.172
Lag 3.CV	0.224	0.262	-0.211	0.295	-0.129	0.157	0.168	0.149	0.116	0.108	0.278	0.177	0.129	0.172	0.032	0.163
Lag 1.DY	-0.161	0.285	-0.034	0.322	-0.008	0.171	-0.301	0.163	-0.090	0.117	0.065	0.194	-0.270	0.188	-0.189	0.178
Lag 2.DY	0.356	0.298	0.243	0.336	-0.095	0.178	-0.147	0.170	0.098	0.122	0.201	0.202	0.142	0.196	-0.315	0.186
Lag 3.DY	-0.413	0.292	0.376	0.329	-0.132	0.175	-0.434	0.166	-0.074	0.120	-0.182	0.198	-0.182	0.192	-0.134	0.182
Lag 1.SH	0.072	0.397	0.392	0.448	0.052	0.238	-0.031	0.226	0.050	0.163	0.001	0.269	0.161	0.261	0.179	0.248
Lag 2.SH	0.343	0.399	-0.177	0.451	0.211	0.240	-0.034	0.228	0.107	0.164	-0.277	0.271	0.146	0.263	-0.059	0.250
Lag 3.SH	0.638	0.414	-0.819	0.467	-0.015	0.249	-0.065	0.236	-0.008	0.170	0.081	0.281	-0.281	0.273	-0.409	0.259
Lag 1.SW	-0.021	0.224	-0.086	0.253	0.452	0.134	-0.011	0.128	0.028	0.092	-0.252	0.152	0.140	0.147	0.297	0.140
Lag 2.SW	0.114	0.250	0.272	0.282	0.312	0.150	0.148	0.143	0.131	0.103	-0.141	0.170	0.046	0.165	0.635	0.156
Lag 3.SW	0.119	0.296	-0.217	0.334	0.265	0.177	0.102	0.169	0.122	0.122	-0.129	0.201	-0.004	0.195	0.284	0.185
Lag 1.WS	0.054	0.263	-0.255	0.297	-0.035	0.158	-0.029	0.150	0.104	0.108	0.324	0.178	-0.126	0.173	-0.238	0.164
Lag 2.WS	0.224	0.240	0.173	0.270	0.049	0.144	0.193	0.137	-0.093	0.098	0.127	0.163	-0.407	0.158	-0.022	0.150
Lag 3.WS	0.001	0.277	0.148	0.313	-0.201	0.166	0.025	0.158	-0.206	0.114	-0.085	0.188	-0.242	0.182	0.064	0.173
Lag 1.WV	-0.554	0.297	-0.138	0.335	0.117	0.178	0.137	0.169	-0.070	0.122	-0.099	0.202	-0.181	0.196	-0.370	0.186
Lag 2.WV	-0.122	0.264	-0.128	0.298	-0.166	0.159	0.116	0.151	-0.187	0.109	-0.233	0.179	-0.115	0.174	-0.223	0.165

Lag 3.WV	0.183	0.298	-0.115	0.337	0.023	0.179	0.076	0.170	-0.210	0.123	-0.102	0.202	0.011	0.196	0.161	0.186
Lag 1.FIVE YEAR FORWARD	-15.095	9.939	11.885	11.211	-5.143	5.962	-4.017	5.665	1.501	4.084	-9.511	6.742	1.045	6.541	-5.272	6.209
Lag 1.MONTHLY BANK RATES	145.84	59.445	-44.884	67.049	76.148	35.660	11.002	33.881	53.952	24.429	11.351	40.323	52.668	39.122	81.486	37.136
Lag 1.Non working Birmingham	-1.176	3.158	2.458	3.562	3.016	1.894	2.692	1.800	-1.422	1.298	-0.816	2.142	1.029	2.078	1.388	1.973
Lag 1.Non working Coventry	0.352	3.762	-3.219	4.243	-3.838	2.257	-4.091	2.144	0.583	1.546	-1.205	2.552	-0.684	2.476	-2.313	2.350
Lag 1.Non working Wolverhampton	-4.057	2.780	-3.868	3.136	1.182	1.668	1.032	1.585	-3.311	1.143	0.359	1.886	-0.765	1.830	0.616	1.737
Lag 1 Precept Stnd Coventry	-31.734	19.708	14.359	22.230	-1.558	11.823	13.151	11.233	3.629	8.099	28.402	13.369	-9.080	12.971	-27.383	12.312
Lag 1 Precept Stnd Birmingham	14.319	14.786	-49.851	16.678	-31.849	8.870	-10.740	8.428	2.853	6.076	-17.989	10.030	-10.285	9.731	-0.760	9.237
Lag 1 Precept Stnd Wolverhampton	20.041	19.046	35.483	21.482	28.391	11.425	-7.853	10.855	-6.387	7.827	-14.582	12.919	19.216	12.534	25.359	11.898
Lag 1 Ctax Stnd Coventry	7.707	5.489	6.057	6.192	9.718	3.293	-2.413	3.129	2.840	2.256	-7.368	3.724	5.940	3.613	5.791	3.429
Lag 1 Ctax Stnd Birmingham	-5.240	4.269	-5.674	4.815	-7.790	2.561	2.198	2.433	-1.766	1.754	5.555	2.896	-4.498	2.810	-3.960	2.667
Lag 1 Ctax Stnd Wolverhampton	3.745	2.624	2.881	2.960	4.341	1.574	-1.096	1.496	0.662	1.078	-3.643	1.780	2.802	1.727	2.641	1.639
Lag 1 School Abs Stnd Coventry	0.297	0.688	-0.525	0.776	-2.333	0.413	-0.412	0.392	-0.334	0.283	-0.918	0.467	0.269	0.453	0.041	0.430
Lag 1 School Abs Stnd Birmingham	2.933	1.513	-2.004	1.706	3.149	0.908	-0.161	0.862	1.239	0.622	-0.826	1.026	0.403	0.996	0.632	0.945
Lag 1 School Abs Stnd Wolverhampton	-1.291	0.774	1.304	0.873	-0.176	0.464	0.034	0.441	-0.279	0.318	1.038	0.525	-0.227	0.509	-0.537	0.483
Lag 1 Free meals Stnd Coventry	-0.838	2.951	1.558	3.328	-7.084	1.770	-0.261	1.682	-2.406	1.213	-1.255	2.002	0.073	1.942	-0.688	1.843
Lag 1 Free meals Stnd Birmingham	32.354	14.887	-18.193	16.792	25.881	8.931	0.530	8.485	7.725	6.118	-13.157	10.098	3.840	9.798	8.549	9.300
Lag 1 Free meals Stnd Wolverhampton	0.757	1.364	-1.994	1.538	-1.169	0.818	-0.738	0.777	0.054	0.560	-1.483	0.925	-1.830	0.897	-2.022	0.852

9.1.3 Terms Used

Term	Definition
ARDL	Auto- regressive Distributed Lag
ARIMA	Auto- regressive Integrated Moving Average Model
Auto-correlation	How today's level of a variable is impacted by lagged levels of the variable
Cross Validation	A resampling approach where a subset of data is used for training and another subset is used for testing.
ETS	Exponential Smoothing
MIDAS	Mixed Data Sampling
Multivariate time series	Allowing a number of variables to interact to forecast
Order of integration	The minimum number of differences required to achieve stationarity
Stationarity	A series is stationary if its properties are not time dependent
Univariate time series	Using a variable as the sole predictor of the levels of that variable
VAR	Vector Auto- regression

10 References

- Dickey, David A., and Wayne A. Fuller. 1979. "Distribution of the Estimators for Autoregressive Time Series with a Unit Root." *Journal of the American Statistical Association* 74 (366). [American Statistical Association, Taylor & Francis, Ltd.]: 427–31. <http://www.jstor.org/stable/2286348>.
- Ellis, Arthur K. 1995. *Teaching and Learning Elementary Social Studies*. ERIC.
- Holt, Charles C. 2004. "Forecasting Seasonals and Trends by Exponentially Weighted Moving Averages." *International Journal of Forecasting* 20 (1): 5–10. <https://doi.org/https://doi.org/10.1016/j.ijforecast.2003.09.015>.
- International Monetary Fund, International Financial Statistics, and data files. 2020. "Inflation, Consumer Prices (Annual)." <https://data.worldbank.org/indicator/FP.CPI.TOTL.ZG?end=2019&locations=GB&start=2010>.
- Kwiatkowski, Denis, Peter C.B. Phillips, Peter Schmidt, and Yongcheol Shin. 1992. "Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root: How Sure Are We That Economic Time Series Have a Unit Root?" *Journal of Econometrics* 54 (1): 159–78. [https://doi.org/https://doi.org/10.1016/0304-4076\(92\)90104-Y](https://doi.org/https://doi.org/10.1016/0304-4076(92)90104-Y).
- "LGInform: The Local Area Benchmarking Tool from the Local Government Association." 2021. <https://lginform.local.gov.uk/>.
- McKendrick, E. 1997. "Local Authorities and Swaps: Undermining the Market?" *Making Commercial Law: Essays in Honour of Roy Goode*.
- "Nomis: Official Labour Market Statistics." 2021. <https://www.nomisweb.co.uk/>.
- Norko, Michael, and Madelon Baranoski. 2008. "The Prediction of Violence; Detection of Dangerousness." *Brief Treatment and Crisis Intervention* 8 (February): 73–91. <https://doi.org/10.1093/brief-treatment/mhm025>.
- Pesaran, M Hashem, Yongcheol Shin, and others. 1995. "An Autoregressive Distributed Lag Modelling Approach to Cointegration Analysis." Department of Applied Economics, University of Cambridge Cambridge.
- Pratt, Travis C., and Francis T. Cullen. 2005. "Assessing Macro-Level Predictors and Theories of Crime: A Meta-Analysis." *Crime and Justice* 32. [University of Chicago Press, University of Chicago]: 373–450. <http://www.jstor.org/stable/3488363>.
- Sims, Christopher A. 1980. "Macroeconomics and Reality." *Econometrica* 48 (1). [Wiley, Econometric Society]: 1–48. <http://www.jstor.org/stable/1912017>.
- Stoddard, Sarah A., Justin E. Heinze, Daniel Ewon Choe, and Marc A. Zimmerman. 2015. "Predicting Violent Behavior: The Role of Violence Exposure and Future Educational Aspirations During Adolescence." *Journal of Adolescence* 44: 191–203. <https://doi.org/https://doi.org/10.1016/j.adolescence.2015.07.017>.
- Taylor, Ralph B., Jerry H. Ratcliffe, and Amber Perenzin. 2015. "Can We Predict Long-Term Community Crime Problems? The Estimation of Ecological Continuity to Model Risk Heterogeneity." *Journal of Research in Crime and Delinquency* 52 (5): 635–57. <https://doi.org/10.1177/0022427815586062>.
- Tong Howell. (2011) Nonlinear Time Series Analysis. In: Lovric M. (eds) International Encyclopedia of Statistical Science. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-04898-2_411
- Widom, Cathy Spatz. 1989. "The Cycle of Violence." *Science* 244 (4901). American Association for the Advancement of Science: 160–66.

———. 1992. *The Cycle of Violence*. US Department of Justice, Office of Justice Programs, National Institute of

Winters, Peter R. 1960. "Forecasting Sales by Exponentially Weighted Moving Averages." *Manage. Sci.* 6 (3). Linthicum, MD, USA: INFORMS: 324–42. <https://doi.org/10.1287/mnsc.6.3.324>.