LUISS

Department of **Political Science**

Degree Program in **Politics: Philosophy and Economics**

Course of **Statistics**

Fiscal Policies and Income Inequality: the Italian Case (2000-2020)

Professor

Giovanni De Luca

104542

Benedetta Bellesia

Academic Year 2024/2025

Index

1. Introduction 3

2. Literature Review and Foundations of Econometric Analysis 6

3. Overview of Econometric Theory 9
- 3.1 About Ordinary Least Squares Models 9
- 3.2 Time Series Analysis and Autoregression Models 14
- 3.3 The Vector Autoregression (VAR) Model 16

4. Descriptive and Inferential Statistics 18
- 4.1 The Dependent Variable and the Selection of Independent Variables 18
- 4.2 Stationarity of Variables 21
- 4.3 The Augmented Dickey-Fuller (ADF) and Bai-Perron Tests 25
- 4.4 Orders of Integration and Cointegration 31

5. Error Correction Model (ECM) and Results 38

6. Limitations and Conclusion 43

7. Bibliography 45

1. **Introduction**

In the recent decades, income inequality has become an increasingly prominent topic in economic and political discourse, with growing concern about its social and economic consequences. With income inequality, we refer to the extent to which income is evenly distributed within a population. There are many factors that can improve or worsen income inequality, both global and domestic. On a global scale, shocks and recessions can send national economies into crises which are difficult to recover from. On a domestic level, fiscal and labour policies can impact the living and working conditions of businesses and individuals, directly influencing wealth accumulation and distribution. This means that income inequality can be affected by both macroeconomic and microeconomic happenings. Understanding how they impact individuals in real life is crucial for designing more equitable and resilient economic systems.

Berg & Ostry (2011) state that a certain level of income inequality is integral in market economies. In fact, income inequality is not to be understood as inherently negative. Inequality can be essential to spur growth and investment, although the authors caution that excessive inequality can become counterproductive, ultimately undermining economic growth and stability. It is thus important to keep inequality levels in check, both to keep the economy stable and to provide equitable living conditions to all. Italy provides a compelling case study for this kind of analysis. In the 21st century, the country experienced many economic fluctuations, including the global financial crisis and the COVID-19 pandemic, as well as several fiscal reforms. These events provide a natural context to examine how tax policy and unemployment interact with income inequality over time.

This thesis attempts to assess whether the employed variables help explain fluctuations in income inequality, both on the long-run and short-run, across a period marked by economic turbulence and policy changes. More specifically, the following research focuses on how taxation and labour market conditions in Italy might have impacted income inequality in Italy between 2000 and 2020 included. The motivation behind this research stems from the highly contested nature of taxation policies themselves. Because their role and effect on the economy is constantly subject of debate, this study aims to analyse whether changes in tax structures, such as adjustments to value-added taxes or wealth taxes, or fluctuating market conditions have effectively contributed to reducing (or exacerbating) income inequality. This is made possible by employing statistical methods and econometric tools. The focus is on time series econometrics, particularly the use of unit root testing, cointegration analysis, and Error Correction Models (ECMs). These tools are employed to distinguish between short-term fluctuations and long-term equilibrium relationships among the variables of interest. It must be addressed that much of the existing literature on the dynamic relationship between macroeconomic indicators relies on vector autoregression (VAR) and vector error correction models (VECMs). These models are well-suited to capturing the interdependencies among multiple non-stationary time series. Although this thesis ultimately employs the Engle-Granger Error Correction Model due to its simplicity and suitability for bivariate analysis, the theory behind the VAR framework is nonetheless presented in this thesis. The VAR model remains highly relevant in understanding the theoretical structure and dynamic interactions that motivate this research.

For this reason, this thesis is structured as follows: the second chapter presents an in-depth literature review on the evolution of Econometrics and applied research in Economics, up until the birth and impact of VAR models; the third chapter, instead, is entirely dedicated to explaining the theoretical basis of this research, specifically the limitations of ordinary regression models, the usefulness of time series analysis in this context, and the properties of autoregression models. Chapters four and five detail the empirical process: first through descriptive statistics and variables selection, then through the appropriate inferential statistics, providing the basis for the final Error Correction Model (ECM). The sixth chapter finally builds the model and interprets the results, while the seventh and last chapter provides a final overview of the entire process and addresses some analytical limitations.

2. Literature Review and Foundations of Econometric Analysis

Econometrics has been fundamental in economics and policy analysis to develop models which can accurately analyse, explain and forecast macroeconomic data. At the end of the 19^{th} century, Francis Galton had just introduced the concept of regression to explain the correlation between height of parents and their children (Galton, 1886). Soon, scholars attempted to apply modern statistical methods to economics: Henry Moore's *"Laws of Wages: An Essay in Statistical Economics"* (1911) was one of the first serious attempts to apply modern statistical methods to economics, by applying regression techniques to analyse the relationship between wages and economic variables like prices and productivity. A little over a decade later, R. A. Fisher developed the core of inferential statistics. Although he was not an economist, his statistical innovations introduced hypothesis testing, sampling distributions, and experimental design principles (Fisher, 1924).

Another major contribution to the birth of modern econometrics was given by Jygve Haavelmo's "*The Probability Approach in Econometrics*" (1944), for which he won the Nobel Prize for Economic Sciences in 1989. Based upon the already-existing theories of probability and statistical inference, Haavelmo attempted to provide a theoretical foundation to the study of interrelations between economic variables by introducing stochastic models for economic relationships (Haavelmo, 1944). However, the conceptual foundations of time series modelling seem to trace back to the Yule-Walker equation developed between the 1920-30s (Yule, 1927; Walker, 1931). This equation showed how current values of a variable can be explained by its own lags, meaning its own past values. This idea formed the basis for the development of Autoregressive (AR) models, such as AR(1) and more generally AR(p) models. These models have been widely appreciated for their simplicity and interpretability. However, their scope is limited: AR models are designed for univariate analysis, making them not well-suited for capturing interactions among multiple variables or for policy analysis, which instead often requires a multivariate approach.

In parallel to the development of time series methods, the post-WWI period saw the emergence of the first large-scale macro-econometric models. In particular, Jan Tinbergen developed one of the first dynamic econometric models based on the Dutch national economy in the 1930s (Assous & Carret, 2022), for which he won the first Nobel Prize for Economic Sciences in history. Tinbergen's work inspired other also Nobel-awarded economists like Lawrence Klein, who developed the standard Keynes' analytical system into a large-scale macro-econometric model for the U.S. national economy (Klein, 1950). The limitations of these models were highlighted in the 1970s by economist Robert Lucas, in his famous "Lucas Critique". In summary, he criticized the existing macro-econometric models for being structurally unstable and failing to account for how economic agents adjust their behaviour in response to policy changes (Lucas, 1976). This critique eventually set the stage for further methodological innovations.

In the meantime, time series analysis saw important advances with the development of Autoregressive Moving Average (ARMA) model (Moran & Whittle, 1951) and the more recent Autoregressive Integrated Moving Average (ARIMA) model (Box, Jenkins, Reinsel, & Ljung, 2015). ARMA models extended simple autoregressive frameworks by incorporating both past values (AR terms) and past forecast errors (MA terms), offering a more flexible structure for modelling stationary time series. However, many economic time series, due to the presence of macroeconomic variables such as GDP, inflation, or tax revenues exhibit non-stationary behaviour. This is a concept that will be introduced and thoroughly explained further. For now, be it sufficient to say that ARIMA models were introduced to account for such non-stationarity.

Further innovations regard the Vector Autoregression (VAR) model, conceptualized during the late 1970s. Building on the precedent autoregressive models, Christopher Sims (1980) proposed a new framework for empirical macroeconomic analysis that offered an innovative approach to modelling dynamic relationships among multiple time series, something which was not possible with the precedent models. Sims argued that up until then, many large-scale macro-econometric models were constructed by relying on imposed strong a priori restrictions. This meant that economists were manually specifying the relations between variables, assuming correlations about the interaction of variables with each other. Sims (1980) challenged this view, arguing that such restrictions could lead to misleading inferences about economic relationships, instead proposing a more flexible, econometric system which would allow data to inform about the dynamic interrelation among variables. From here onwards, macro-econometric frameworks have undergone important theoretical and methodological advancements, with researchers constantly seeking to improve their interpretability and policy relevance.

VAR models have been extensively used in the analysis of the dynamic effects of fiscal policy on macroeconomic variables. Some pioneering work has been done by Blanchard & Perotti (2002) on the dynamic effects of shocks in government spending and taxes on U. S. activity in the postwar period. In fact, many economic recovery policies have been rooted in the Keynesian perspective on fiscal intervention, which emphasizes the positive role of government spending in stimulating aggregate demand during downturns. This Keynesian consensus was challenged in the subsequent decades, leading scholars to seek empirical tools to improve macroeconomic forecasting. The results were insightful from a policy-making perspective and they underscore the usefulness of empirical and econometric approaches in conducting in-depth economic analyses: the authors found that an increase in government spending corresponded to an increase in private consumption, which supports Keynesian economics, but also to a (slight) decrease in private investment, which instead supports neoclassical theories (Blanchard & Perotti, 2002). Although this may seem like a paradox, it reflects the complex and often mixed nature of fiscal policy transmission mechanisms. This dual outcome highlights the limitations of applying any single theoretical framework universally and reinforces the need for empirical models to capture such nuanced economic dynamics.

Empirical models are also able to compute global forecasting, as shown by Ricci-Risquete and Ramajo-Hernández (2014). In their paper, the authors estimate a Global VAR (GVAR) model for fourteen countries of the former EU15 and the United States (USA), using quarterly macroeconomic, monetary and fiscal data from 1978 to 2009. It only takes a brief calculation to understand that, in total, there are 128 observations per variable (32 years times 4 quarters). The results are particularly important from a macroeconomic policy perspective, as the authors were able to analyse the transmission of fiscal shocks not only within national borders but also across countries, accounting for global interdependencies. Similar to Blanchard and Perotti (2002), their findings highlight a heterogeneous response of private consumption and investment to fiscal policy shocks: indeed, while tax increases generally had a contractionary effect on private consumption, government spending had mixed effects depending on the country and the nature of the shock. Moreover, the authors found similarities in the cyclical behaviour of European economies, thus recommending a coordination of European fiscal initiatives to counterattack unwanted spillovers of domestic shocks. This conclusion highlights the usefulness of econometric models not just for national economies, but also regional ones, if not global.

3. Overview of Econometric Theory

Linear regression models have long served as the foundational tool in empirical economics, due to its linear structure and interpretability, which make it attractive for early-stage analyses and policy evaluations. However, when the focus shifts towards understanding how economic policies unfold over time, especially in macroeconomic contexts such as fiscal interventions, linear regressions begin to show their limits. Time-dependent variables, structural breaks, policy shifts, and feedback mechanisms often violate key assumptions of OLS regression models, leading to biased, inconsistent, or misleading estimates.

This chapter explores these limitations by grounding the discussion in the research question: "How have Italian fiscal policies implemented between 2000 and 2020 impacted income inequality in Italy?" Using this context, it illustrates how standard linear regression frameworks struggle with multidimensional policy instruments and temporally correlated data. It introduces the concept of omitted variable bias, highlights the difficulty of isolating the effect of one policy in a complex policy environment, and demonstrates how key OLS assumptions are routinely violated when dealing with time series data. The inadequacy of OLS in this context sets the stage for the subsequent chapters, in which the empirical analysis will be conducted along with an explanation of the appropriate, chosen econometric model. The entirety of the econometric theory illustrated in this chapter and in the subsequent sections is derived from course materials and the mandatory textbook *Statistics: Principles and Methods* by Cicchitelli, P., D'Urso, M., & Minozzo, M. (2021), integrated with the academic textbooks *Introduction to Econometrics* (4th ed.) by Stock, J.H., & Watson, M.W. (2020) and *Applied Econometric Time Series* (4th ed.) by Enders, W. (2014).

3.1 About Ordinary Least Squares Models

To understand the limitations of OLS regression, it helps to first look at how a basic OLS model works. In its simplest form, OLS regression analysis is used to examine how one variable affects another. For instance, we might want to see how a change in tax rates (our explanatory variable) influences income inequality (our outcome variable). The relationship between these two variables can be represented by a simple linear equation:

$$Y = \beta_0 + \beta_1 X + \varepsilon,$$

where X is the explanatory variable, called independent variable; Y is the outcome variable, called dependent variable; β_0 is one of two parameters, called the intercept, and it is equal to the expected value of Y when $X = 0$; β_1 is the second parameter, it represents the slope of the regression line. Finally, ε is called the error term.

Ordinary Least Squares (OLS) is a statistical method used to estimate the relationship between one (or more) independent predictors and a dependent variable. In the presence of just one predictor, the relationship is estimated by finding the line that minimizes the sum of the squared differences (called residuals) between the observed values and the predicted values. This "least squares" approach ensures the best linear fit to the data, and it is visualized as the function of a straight line.

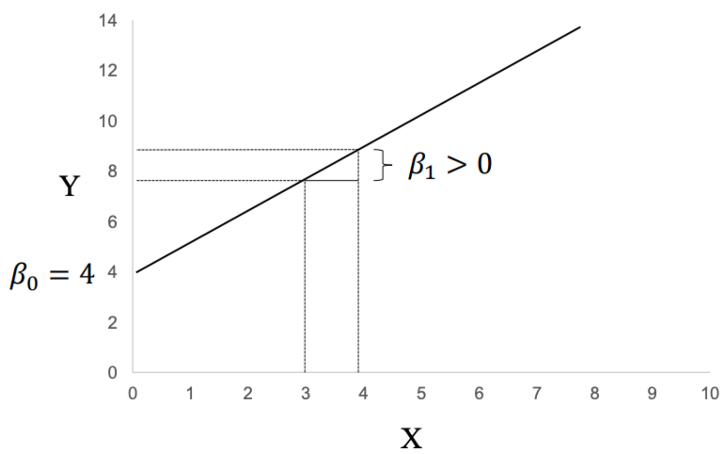

Figure 1. Simple Linear Regression Model (Source: Course material)

For the time being, we will leave the meaning of "error term" aside to focus on explaining the relationship between dependent and independent variables. In any empirical analysis, understanding the relationship between variables is fundamental. Specifically, researchers often aim to assess how one or more factors influence an outcome of interest. The dependent variable is the element being explained or predicted, while the independent variables are those believed to have an effect on it. In the context of this research, the independent variable X would be represented by the tax rates, while the outcome variable Y would be represented by income inequality.

This is when some problems start showing: first, we are faced with a conceptual problem: in fact, "tax policies" is a concept too broad and multidimensional to be represented by a single variable, a condition referred to as conceptual stretching. When a concept is applied too broadly or imprecisely, it eventually loses its analytical usefulness or meaning. In economic research, this is quite problematic, as a term as broad as "tax policies" encompasses a range of distinct instruments such as income tax rates, corporate taxation, value-added tax, and wealth taxes. Each of these can exert different effects on income inequality, so that even if we did decide to focus on one specific instrument, we would necessarily omit other instruments which equally influence our outcome variable Y. This kind of proceeding would then completely miscalculate the statistical model, leading to misleading results. Thus, in the context of this specific research, a single-variable model may fail to account for the various mechanisms at play and may risk excluding important explanatory variables.

It is exactly this realization which leads us to a second, fundamental issue, for which we need to introduce the concept of "error term". We can then define the error term ε as everything that affects the dependent variable Y but is not included in the model. In other words, it is the difference between the actual value and the value that the true model (which we never fully know) would predict. In practice, we use the residuals as estimates of these errors: the smaller the residuals, the better the model fits the data. Indeed, analysing the impact of one single, broad predictor leaves out a number of omitted variables that can notably alter the statistical model. If such predictors are not accurately represented in the regression model, they are included in the errors, creating omitted variable bias.

For there to be omitted variable bias, two conditions must be satisfied: first, the omitted variable must be related to the independent variable X; second, the omitted variable must also affect the dependent variable Y. Essentially, we find omitted variable bias when the omitted variable affects both the explanatory and the outcome variable. This is relevant for the sake of the present analysis: by identifying X with a single tax policy, such as the Italian income and personal tax (IRPEF), multiple other factors are ignored. These factors become our missing variables. It is not simply about other taxation policies, such as the

Italian value added tax (IVA). In fact, macroeconomic analyses encompass variables as inflation rates, social spending, employment rates, etc… which not only influence the outcome variable *Y*, but also the explanatory variable *X* itself. Omitting them from the model leads to biased and potentially misleading regression estimates.

Omitted variables can therefore make the OLS estimators misleading or biased. Unfortunately, this issue cannot be solved by simply employing a multiple regression analysis where instead of having one regressor, we make use of multiple variables X_i. To understand why, let's build a hypothetical multiple linear regression model. Assume our variables are the main financial instruments that affects individual income in Italy: the personal income tax (Imposta sul Reddito delle Persone Fisiche, IRPEF), the corporate income tax (Imposta sul Reddito delle Società, IRES), the value-added tax (Imposta sul Valore Aggiunto, IVA), the municipal property tax (Imposta Municipale Unica, IMU), and the tax on gifts and inheritance (Imposta sulle Donazioni e Successioni); additionally, government expenditures on social welfare, education, and healthcare, as well as key macroeconomic indicators including the inflation rate, unemployment rate, and GDP growth. Our function should the look like as follows:

$$Y = \beta_0 + \beta_1\text{IRPEF} + \beta_2\text{IRES} + \beta_3\text{IVA} + \beta_4\text{IMU} + \beta_5\text{Gift\_Tax} + \beta_6\text{SocialSpending} + \beta_7\text{EducationSpending} + \beta_8\text{HealthcareSpending} + \beta_9\text{Inflation} + \beta_{10}\text{Unemployment} + \beta_{11}\text{GDP\_Growth} + \varepsilon$$

Although the model now seems to be more responsive and accurate, by checking its compliance with the OLS assumptions we would soon come to the realisation that this is not the case. These assumptions form the foundation of OLS regression: they ensure that the model produces unbiased, consistent, and reliable estimates of the relationship between variables. If one or more of these assumptions is violated, the results may be misleading, even if the model appears statistically sound on the surface.

The first OLS assumption states that large outliers, meaning observations with values far outside the usual range of the data, are unlikely; second, the data must be drawn from the population through a random sampling process, meaning each observation is independently and identically distributed (i.i.d.); third, the data must not present perfect multicollinearity, meaning that there should be no exact linear relationship between the independent variables. Finally, the fourth OLS assumption states that $E(\varepsilon_i \mid X_i) = 0$, meaning that the error term has an expected value of zero given any value of the independent variables.

Unfortunately, as a consequence of working with macroeconomic data, the second assumption is automatically violated. In fact, the variables employed are not independent, nor randomly assigned: not only fiscal policies are systematically determined based on economic conditions rather than being exogenous (randomly assigned), but they are also persistent across time. Tax rates, government spending, and social programs do not change randomly each year but evolve based on past policies and economic trends. Furthermore, we are once again faced with omitted variable bias, which directly violates the fourth OLS assumption. In less technical terms, the assumption states that the error term (ε_i) should be uncorrelated with the independent variable (X_i). When analysing macroeconomic variables across time we are often faced with trends, cycles, or structural breaks, which violate the assumption that the error term has constant variance and zero correlation with the regressors. In particular, if key variables are omitted or if the model does not include feedback effects, the error term may become correlated with the independent variables, violating the assumption $E(\varepsilon_i \mid X_i) = 0$. This leads to omitted variable bias and inconsistent estimates, even in large samples. Moreover, OLS assumes no perfect multicollinearity, when one independent variable in a regression model can be exactly predicted by one or more of the other independent variables. In other words, there is a perfect linear relationship between two or more variables in the model. Although macroeconomic data do not necessarily present such characteristic, they do often trend together over time, leading to high multicollinearity, which makes it difficult to isolate the effect of individual predictors. Finally, time series often violate the assumption of stationarity, that is, the statistical properties of the variables (mean, variance, and covariance) remain constant over time. If the data are non-stationary, regression results may be spurious, meaning they appear statistically significant even when there is no real relationship.

Classical linear regressions thus present a major limitation that was only implicitly referred to up until now: the inability to account for time. This means that since linear regressions typically estimate relationships at a single point in time, they ignore past values of variables that may affect current outcomes. Often times, when dealing with macroeconomic variables, the current value of X_i is highly related to its past values. This happens when dealing with time series: a sequence of data points collected or recorded at regular intervals over time. Time series data focus on how a particular variable evolves. Each observation in a time series is ordered chronologically, and the timing of each data point is a crucial aspect of the analysis. Because these values are observed over time, they often exhibit patterns such as trends, seasonality, or cycles, and are influenced by past events, all of which linear regression is unable to account for. By leaving out variables that drive policy changes, the relationships between variables might be estimated wrongly, and the effects misattributed. The result is then a misleading interpretation of the effect of X_i on Y_i.

In conclusion, while multiple OLS regression models offer a clear and intuitive framework for estimating relationships between variables, their assumptions often do not hold when applied to time series data. Therefore, although more complex, time series analysis offers a more appropriate methodological approach for studying data that unfolds over time. The next chapter will introduce autoregressive models, which are specifically designed to address the limitations discussed here and provide a more accurate understanding of economic relationships in a time-dependent context.

3.2 Time Series Analysis and Autoregression Models

Economists have long sought to account for the role of time and endogeneity in statistical analysis, particularly when dealing with macroeconomic data. As shown in the previous section, simple and multiple regression models face significant limitations in this context. They struggle to address issues such as omitted variable bias, autocorrelation, and the temporal structure of economic variables that evolve through feedback mechanisms and policy responses. These limitations make regression analysis poorly equipped to capture the true dynamics of macroeconomic relationships, especially when variables are interdependent and change over time. In response to these shortcomings, time series analysis has emerged as a more appropriate methodological framework. By explicitly incorporating the temporal dimension and allowing for lagged effects, structural shifts, and mutual interdependencies, time series models provide a more robust approach to analysing economic phenomena across time.

Thus, let's assume a given data set where Y is our variable and Y_t is the observation of given variable at date t. Let T be the total number of observations taken into consideration. Finally, let Y_{t-1} also be the first lagged value, that is, the value of Y in the previous period $t - 1$ relative to period t. In time series data, the value of Y_t and its lagged value Y_{t-1} are typically correlated. In econometrics, this correlation is defined as autocorrelation. When analysing economic phenomena that evolve over time, such as regional income disparities and fiscal policy interventions, it becomes essential to account for past values influencing present outcomes. The simplest way to capture this temporal dependence is through a first-order autoregression or AR(1), where the current value of a variable is modelled as a function of its own lagged value. This model can be written in form of a regression model as follows:

$$Y_t = \beta_0 + \beta_1 Y_{t-1} + \varepsilon_{t,},$$

with the variable Y being on both sides of the equation, though expressed differently. In fact, the regressor Y_{t-1} is just the first lag of Y_t. To put it visually, the graph of a AR(1) model with function $Y_t = 18 - 0.8Y_{t-1} + \varepsilon_t$ looks like as shown in Figure 2.

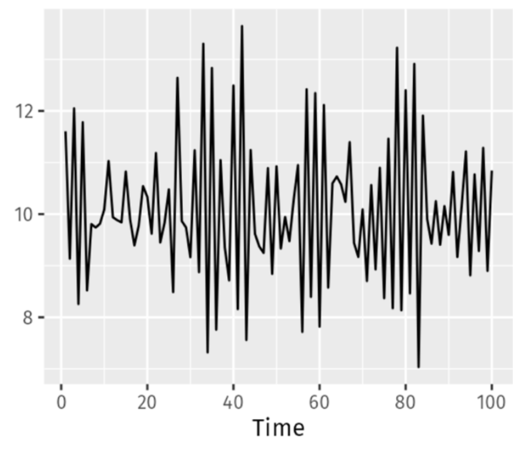

Figure 2. Time Series Plot (Source: 8.3 Autoregressive Models | Forecasting: Principles and Practice (2nd Ed), n.d., 2025)

The AR(1) model uses only one lag to forecast the future. However, doing so risks ignoring potentially useful information in the more distant past. Thus, to include multiple lags, the p^{th}-order autoregressive model AR(p) is preferred. The regression equation looks as follows:

$$Y_t = \beta_0 + \beta_1 Y_{t-1} + \beta_2 Y_{t-2} + \ldots + \beta_p Y_{t-p} + \varepsilon_t,$$

where $\beta_0, \beta_1, \ldots, \beta_p$ are the unknown coefficients, and the number of the lags p is called the order, or the lag length, of the model.

Finally, adding other variables and their lags to an autoregression can improve forecasting performance. In fact, macroeconomic systems are rarely univariate; they are composed of multiple interrelated indicators that move together over time. For instance, government spending in one region may not only be influenced by its own past trends but also by changes in regional income, unemployment, or economic shocks. Thus, AR models use past values of a single variable to predict its current or future values. One way to include several variables in a single model is to use a Vector Autoregression (VAR).

3.3 The Vector Autoregression (VAR) Model

To address the problem introduced earlier, this chapter introduces the vector autoregression (VAR) model. Unlike linear regression models, the VAR framework accounts for endogeneity by treating all variables as jointly endogenous and explicitly modelling their lagged interdependencies. This means that each variable is modelled not only as a function of its own past values, but also as a function of the past values of all other variables in the system. In doing so, the VAR framework captures the dynamic interdependencies and feedback mechanisms that are often present in macroeconomic relationships. As a result, the VAR approach is well suited for analysing dynamic relationships in macroeconomic time series data, where strict OLS assumptions may not hold. In the case of two time series variables X and Y, the model equation looks as follows:

$$Y_t = \beta_{10} + \beta_{11} Y_{t-1} + \ldots + \beta_0 + \beta_{1p} Y_{t-p} + \gamma_{11} X_{t-1} + \ldots + \gamma_{1p} X_{t-p} + \varepsilon_{1t},$$
$$X_t = \beta_{20} + \beta_{21} Y_{t-1} + \ldots + \beta_0 + \beta_{2p} Y_{t-p} + \gamma_{21} X_{t-1} + \ldots + \gamma_{2p} X_{t-p} + \varepsilon_{2t}.$$

In the equations, β's and γ's are unknown coefficients, while u_{1t} and u_{2t} are error terms. The VAR model is useful as it extends the univariate autoregression to multiple time series variables, making it a multivariate autoregression. Thus, a vector autoregression is just made of k regression equations for k number of variables.

In essence, the VAR model generalizes the concept of univariate autoregression to a multivariate setting, allowing for a richer representation of the data-generating process. This is especially useful for policy-related research, which investigates the dynamic interaction of multiple variables at once. By incorporating past dynamics and feedback effects, the VAR framework offers a more realistic and informative analysis than static or univariate approaches.

Because VAR models are essentially made of multiple equations, inference in VARs can also be estimated according to OLS assumptions: first, all variables included in the model must be stationary; second, the error term must once again have conditional mean of zero; third, large outliers are unlikely; finally, there must be no perfect collinearity. This may seem like a paradox, as these resemble the classical OLS assumptions which we refuted earlier. Yet, VAR models can be estimated using Ordinary Least Squares (OLS) because even though the full VAR model is multivariate, each individual equation can be estimated separately, provided that certain OLS assumptions still hold. The main difference is that we do not assume exogeneity, meaning we do not assume that a single variable is independent from other variables in the model.

One great limitation of VAR models is that they are inherently atheoretical: they do not impose any prior economic theory or structural assumptions on the relationships between variables. Instead, they and purely data-driven, which means that they let data speak for themselves. While this flexibility may represent a strength, it also requires cautious interpretation.

4. Descriptive and Inferential Statistics

This chapter outlines the econometric strategy employed to investigate the relationship between fiscal policy instruments and income inequality in Italy. It begins by describing the data sources and variable selection, followed by the steps taken to ensure the validity of the time series analysis, including stationarity tests and model specification. The core of the chapter presents the estimation of a Vector Autoregression (VAR) model, which captures the dynamic interdependencies among the selected variables. Finally, the results are interpreted through impulse response functions and variance decomposition to assess the impact of fiscal measures over time.

4.1 The Dependent Variable and the Selection of Independent Variables

The dependent variable employed in this research is income inequality, measured as the Gini coefficient. It is one of the most internationally recognised and widely employed indicators of inequality, as it provides a single, standardized value that reflects the overall distribution of income within a country. The values of the Gini coefficient scale from a minimum of 0 to a maximum of 1. The higher the Gini coefficient, the higher the level of income inequality in a country. Its standardized value makes it comparable across time and different nations, which was deemed fitting for the sake of the present analysis.

Then, several potential independent variables were identified as predictors of income inequality in Italy. Firstly, there are the key fiscal indicators: the tax on personal income and individuals (IRPEF), the corporate income tax (IRES), the value added tax (IVA), the property tax (IMU), and the inheritance and gift tax. Moreover, there are other macroeconomic variables that may influence income inequality in Italy, such as the level of social, education and healthcare spending, as well as the inflation rate, the unemployment rate, and GDP growth rate per capita.

All the observations employed in this analysis, for both dependent and independent variables, were retrieved from major databases like ISTAT, the OECD, Eurostat, and the World Bank. However, it is necessary to specify that the observations for the Gini coefficient, sourced from the World Bank, presented a missing observation for the year 2001. To address this issue, the value for 2001 was imputed using the arithmetic mean of the adjacent years, 2000 and 2002. This method was deemed appropriate, given the relatively stable year-over-year values of the income inequality measure.

This study investigates the relationship between fiscal policy and income inequality in Italy over a defined period of time. Specifically, it focuses on the years 2000 to 2020, providing a total of 21 annual observations. This time frame was chosen based on the availability and consistency of macroeconomic data from national and international statistical sources. In fact, while a single missing data point can reasonably be addressed through methods such as calculating the arithmetic mean, the presence of multiple missing values across different variables represents a greater challenge. As a result, the period prior to 2000 was excluded from the analysis to preserve data consistency and model reliability.

The analysis will thus be based on time series data that only covers 21 years. Thus, it is important to limit the number of independent variables in the model. This is because each additional variable reduces the model's degrees of freedom, meaning we have fewer data points left to estimate the effect of each variable accurately. With too many independent variables and too little data, the model becomes unstable, and the results can be unreliable or misleading. This is known as overfitting: the model tries too hard to match the specific data points, rather than identifying general patterns. In short, with a small sample size, using fewer well-chosen variables helps ensure that the results are more robust and meaningful.

To reduce the number of predictors while preserving explanatory power, a Pearson correlation matrix needs to be constructed to assess the strength of the linear relationships between each independent variable and the dependent variable. Thankfully, Excel has built-in tools that allow users to perform basic descriptive analyses. Thus, a matrix of correlation with all variables was built, and values exhibiting a correlation inferior to |0.5| with the dependent variable were excluded from further analysis.

Variable	IRPEF	IRES	IVA	IMU	Gift Tax	Social Spending	Education Spending	Healthcare Spending	Inflation Rate	Unemployment Rate	GDP Growth	Income Inequality
IRPEF	1.000											
IRES	-0.452	1.000										
IVA	0.098	-0.316	1.000									
IMU	0.489	-0.567	0.403	1.000								
Gift Tax	-0.004	0.050	0.561	0.190	1.000							
Social Spending	0.799	-0.763	0.240	0.657	-0.049	1.000						
Education Spending	-0.413	0.535	-0.404	-0.615	-0.063	-0.416	1.000					
Healthcare Spending	-0.463	0.571	-0.274	-0.689	-0.013	-0.752	0.311	1.000				
Inflation Rate	-0.392	0.607	-0.254	-0.603	-0.070	-0.727	0.356	0.636	1.000			
Unemployment Rate	0.436	-0.548	0.438	0.842	0.517	0.540	-0.593	-0.481	-0.548	1.000		
GDP Growth	-0.569	0.173	0.471	-0.170	0.297	-0.584	-0.192	0.349	0.239	-0.003	1.000	
Income Inequality	0.218	-0.478	0.543	0.530	0.636	0.416	-0.288	-0.397	-0.248	0.751	0.065	1

Table 1. Correlation Matrix of Socioeconomic and Fiscal Variables (Source: Own Elaboration)

As shown in the picture, the relevant results were as follows: property tax (IMU), r = 0.530; value-added tax (IVA), r = 0.543; unemployment rate, r = 0.751; and gift tax, r = 0.636. A significant issue of multicollinearity was identified between the IMU and unemployment rate variables, with a correlation coefficient of r = 0.842. Thus, IMU was excluded from the final model to mitigate multicollinearity and improve model stability. Slight correlation was found between gift tax with both unemployment rate and IVA, but since the values were moderate, they were found to be negligible. As a result, the regression analysis proceeded using IVA, gift tax, and the unemployment rate as the final independent variables.

Finally, Excel can be a useful and powerful ally when computing basic statistical calculations and basic data organization and visualization. However, it lacks the built-in statistical tools which are essential in time series analysis to determine stationarity, for instance, which is a key requirement. Statistical programming environments like R and RStudio, instead, are more suited for data analysis and econometrics, as they can handle serial autocorrelation and are able to support more advanced statistical models like VAR. From now on, RStudio will be employed due to its user-friendly dashboard and overall greater accessibility.

4.2 Stationarity of Variables

Another concept which is highly relevant in time series analysis is stationarity, which was not introduced earlier for reasons of practicality and simplicity. In OLS regression models, it is assumed that the properties of the variables, such as mean, variance and covariance, are constant over time, meaning they do not depend on the time at which the series is observed. Thus, a time series Y_t is said to be stationary if the joint distribution of (Y_{s+1}, Y_{s+2}, …, Y_{s+T}) does not depend on s regardless of T. If this holds true, the variables are said to be stationary. In other words, the underlying data-generating process should remain stable, implying that the future behaves statistically like the past. However, due to their macroeconomic nature, indicators such as income levels, government spending, or GDP, are often found to be non-stationary, often exhibiting trends, seasonality, or structural breaks. In such cases, the series may display values that are highly correlated over time simply due to underlying trends rather than genuine economic relationships.

Therefore, stationarity is essential to ensure that observed relationships in the data reflect genuine economic associations rather than coincidental movements.

There can be different types of non-stationarity, such as trends and breaks. This persistence can give rise to spurious regression, where two unrelated non-stationary variables appear to be significantly correlated due to shared trending behaviour. Stationarity, thus, is the concept that historical relationships can be generalized to the future: if the variables are said to be stationary, it can be said that values of the future will be similar to values of the past. If this assumption does not hold, the values are said to be nonstationary, which means that the historical data is not reliable to predict the future. It is for this reason that formal stationarity tests should be conducted, and appropriate transformations may be required to achieve stationarity before proceeding with time series analysis. In the context of macroeconomic analysis, it is often reasonable to expect that certain variables may exhibit non-stationarity due to underlying economic trends, structural changes, or policy shifts. Unlike short-term fluctuations, such large-scale economic disruptions tend to have long-lasting effects. This persistence over time may introduce a stochastic trend, meaning the series does not return to a fixed mean and is instead affected by shocks that have permanent impacts.

To test whether our series present non-stationarity, RStudio allows users to employ the Autocorrelation Function (ACF), which measures the correlation between a time series and its own past values at different lags. In other words, it helps identify whether past values of a variable have a systematic influence on its current values.

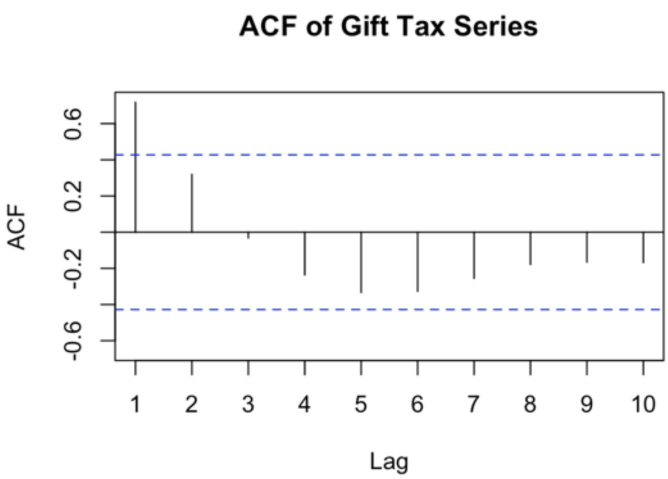

Table 2. (Source: Own Elaboration)

As Table 2 shows, the autocorrelation function (ACF) plot shows a vertical bar for each lag, up to 10 on the X-axis, and a blue dotted line. The height of each bar represents the correlation between the time series and a lagged version of itself at that specific lag. The blue line represents the confidence interval (usually set at 95%) which, if crossed, indicates a statistically relevant autocorrelation. In this case, the ACF plot of the gift tax series reveals a prominent spike at lag 1, which crosses the 95% confidence bounds. This alone is not enough to determine the presence of non-stationarity, although a sudden spike might suggest a structural break. A somewhat similar pattern was observed for the time series of the value-added tax (IVA) in Table 3, where autocorrelation is present at lag 1.

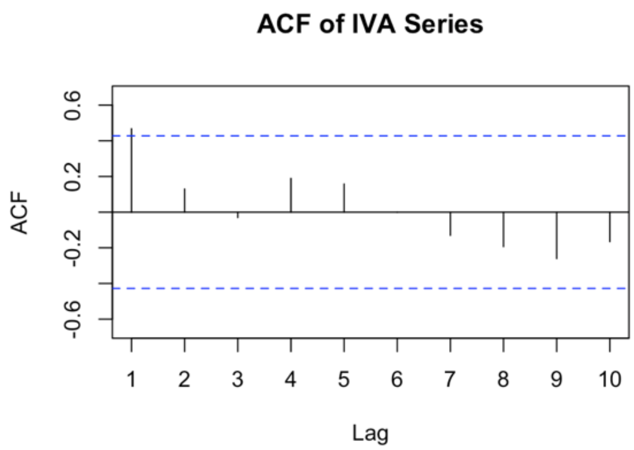

Table 3. (Source: Own Elaboration)

Finally, the unemployment rate series in Table 4 exhibited a strong autocorrelation at lag 1, followed by a gradual decline, which is a typical sign of non-stationarity often found in macroeconomic indicators. This pattern suggests the presence of non-stationarity:

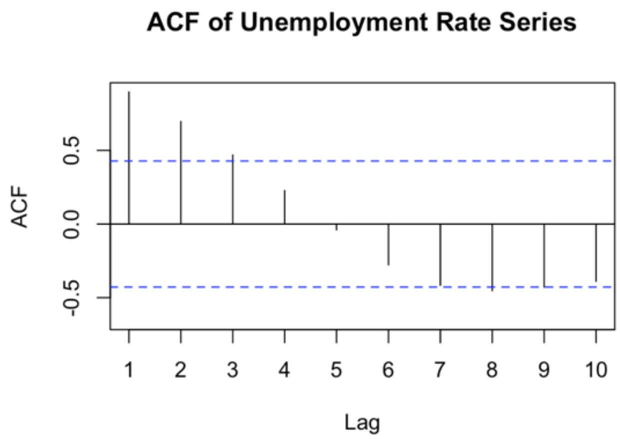

Table 4. (Source: Own Elaboration)

To visualize these spikes and trends in the three time series, we use the *plot()* function:

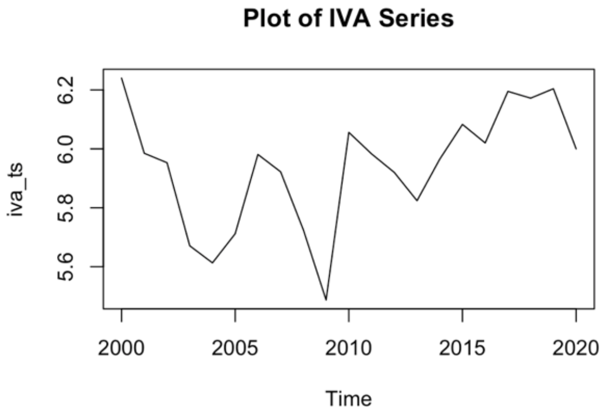

Table 5. (Source: Own Elaboration)

In Table 5, the graph of the IVA series shows a decline in observation 10, which corresponds to the year 2009. In the data retrieved from ISTAT, the values for IVA observations dropped from 5.725 in 2008, to 5.487 in 2009. There indeed have been policy changes between 2000 and 2020, which, however, only periodically increased the value-added tax of a 1%. This information alone is not sufficient to determine with certainty the presence of a complete structural break.

Table 6. (Source: Own Elaboration)

Interestingly, the plot for the gift tax series shows a drop leading to observation 5 and a slight increase right after, corresponding exactly with the evolution of the gift tax in Italy: abolished in 2001, it was reintroduced in 2006 and was kept constant up until the end of our time series. This structural break seems to be confirmed both historically and graphically, although these observations are still not enough to determine with certainty a structural break.

Finally, the plot of the unemployment rate series:

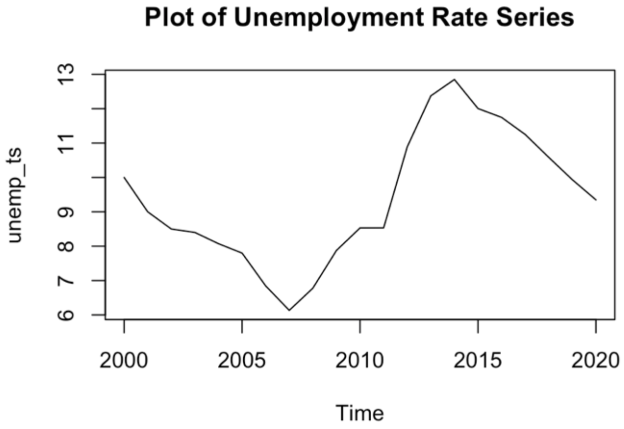

Table 7. (Source: Own Elaboration)

As Table 7 shows, there is a decline leading up to 2007. Then, the line shows a sudden and rapid increase, reaching its peak in 2014. This trend is confirmed historically: according to the Italian Presidency of the Council of Ministers, specifically the *Dipartimento per la programmazione e il coordinamento della politica economica,* in 2007, the unemployment rate in Italy reached its lowest value of 6.13. The crisis fully hit the Italian economy in 2009, when industrial production collapsed, and small-to-medium enterprises (SMEs) struggled, with unemployment rising to a 7.88%. In 2010-2011 the eurozone faced the sovereign debt crisis and in 2012-13 the austerity measures were introduced (Basso, 2016). Although necessary to save the eurozone, they further exacerbated the situation in Italy, with unemployment reaching its peak at nearly 13% in 2014. From there onwards, the peninsula made a slow recovery due to pre-existing structural, bureaucratic and financial issues (Schilirò, 2014).

4.3 The Augmented Dickey-Fuller (ADF) and Bai-Perron Tests

The Augmented Dickey-Fuller (ADF) test is a widely used statistical test to assess the stationarity of a time series, an essential concept in time series analysis. As explained previously, we define "stationary" those time series whose properties, such as mean and variance, do not change over time. This assumption is fundamental in many econometric models, as non-stationary data can lead to unreliable and spurious results. The ADF test is an extension of the Dickey-Fuller test, which tests for the presence of a unit root in a series. In the ADF test, we consider two hypotheses: the null one, which states that the variable is non-stationary. The alternative hypothesis, conversely, suggests that the series is stationary. In order to choose between the two, just like we would do in a regular hypothesis testing, we check the p-value. If $p > 0.05$, then the null hypothesis cannot be rejected, and the series is non-stationary. Thankfully, RStudio allows users to easily compute an ADF test through the *adf.test()* function:

Before computing the ADF test, we need to choose the number of lags to implement in the analysis. Choosing the correct number of lags is paramount: too few lags may leave autocorrelation in the residuals, leading to biased test statistics, while too many lags may overfit the model and reduce the power of the test, especially when the sample size is small. To make this choice, we are free to choose between four information criteria (AIC, HQ, SC/BIC, and FPE) which help with determining the number of lags to use. To do so, the *VARselect()* function from the vars package in RStudio is employed, which tests lag orders up to 10.

When the function is applied to the present analysis, the four selection criteria unanimously indicate that a lag order of 3 is optimal. However, when computing the *adf.test()* function, RStudio immediately warns us that the time series analysed are too short relative to the number of lags (*k* in the function). One possible solution for the issue is to run the same test with $k = 2$ and see if it works. Although this may represent a limitation for the study, it is not an arbitrary decision, as it is imposed by data limitations and practical constraints.

Test	Value
Test Type	Augmented Dickey-Fuller
Variable	iva
Test Statistic	-4.7014
Lag Order	2
P-Value	< 0.01
Alternative	Stationary

Table 8. ADF Test for IVA Series (Source: Own Elaboration)

As Shown in Table 8, when $k = 2$, the time series results to be stationary, with a p-value much lower than 0.05. This confirms our preliminary analysis about the policy changes in the value added tax.

Test	Value
Test Type	Augmented Dickey-Fuller
Variable	gift
Test Statistic	-1.7602
Lag Order	2
P-Value	0.6637
Alternative	Stationary

Table 9. ADF Test for Gift Tax Series (Source: Own Elaboration)

Table 9 shows the ADF test for the gift tax series. As shown, the series' p-value is 0.664 when $k = 2$. This is the complete opposite compared to the IVA series. Given the conflicting results, it is necessary to investigate whether there are structural breaks affecting the outcome. Historical and institutional knowledge about Italy's gift tax policy supports this thesis: the tax was in place in 2000, abolished in 2001, and reintroduced in 2006 at a stable rate, remaining constant until 2020. This pattern suggests the presence of regime changes rather than a genuine trend or stochastic non-stationarity.

Since the ADF test does not account for such structural shifts and may therefore mistake them for non-stationarity, the gift tax and value-added tax (IVA) series were tested employing the Bai-Perron multiple structural break test, which is specifically designed to detect breaks in the mean of a time series. In Table 10, m refers to the breakpoints identified by the test, while the measures of fit on the bottom are the information criteria which were introduced earlier.

The best BIC (identified at its lowest value) for the gift tax series is shown in Table 10 and it is when $m = 3$, which means that structural breaks were identified at observations 3, 8 and 12. Indeed, the test identified breakpoints that correspond closely to the known policy changes: observation 3 refers to the year 2002; observation 8 refers to 2007, while observation 12 refers to 2011, which might be due to minor fluctuations. Overall, the results confirm that the gift tax series exhibits meaningful structural breaks.

Finally, the Bai-Perron test of the IVA time series (Table 11) revealed that while the ADF test regarded the IVA series to be stationary, the Bai-Perron test revealed the presence of possible structural breaks when $m = 2$. This indicates that while the series is stationary overall, its statistical properties shift significantly over time.

Bai-Perron Breakpoints – Gift Tax Series

Segment	Breakpoints	Breakdates
m = 1	3	0.142857142857143
m = 2	3, 9	0.142857142857143, 0.428571428571429
m = 3	3, 8, 12	0.142857142857143, 0.380952380952381, 0.571428571428571
m = 4	3, 8, 12, 15	0.142857142857143, 0.380952380952381, 0.571428571428571, 0.714285714285714
m = 5	3, 8, 12, 15, 18	0.142857142857143, 0.380952380952381, 0.571428571428571, 0.714285714285714, 0.857142857142857

Fit Statistics (RSS & BIC)

Statistic	m_0	m_1	m_2	m_3	m_4	m_5
RSS	8.080e-3	4.043e-3	1.543e-3	1.191e-3	1.171e-3	1.123e-3
BIC	-9.941e+01	-1.079e+02	-1.220e+02	-1.214e+02	-1.156e+02	-1.104e+02

Table 10. Bai-Perron Test for Gift Tax Series (Source: Own Elaboration)

Bai-Perron Breakpoints – IVA Series

Segment	Breakpoints	Breakdates
m = 1	15	0.714285714285714
m = 2	3, 10	0.142857142857143, 0.476190476190476
m = 3	3, 10, 15	0.142857142857143, 0.476190476190476, 0.714285714285714
m = 4	3, 6, 10, 15	0.142857142857143, 0.285714285714286, 0.476190476190476, 0.714285714285714
m = 5	3, 6, 10, 14, 17	0.142857142857143, 0.285714285714286, 0.476190476190476, 0.666666666666667, 0.80952380952381

Fit Statistics (RSS & BIC)

Statistic	m_0	m_1	m_2	m_3	m_4	m_5
RSS	0.8256	0.5722	0.3683	0.2960	0.2740	0.2676
BIC	-2.2739	-3.8864	-7.0509	-5.5464	-1.0830	4.5089

Table 11. Bai-Perron Test for IVA Time Series (Source: Own Elaboration)

In summary, the analysis confirms that our time series require careful treatment before being included in a VAR model. The unemployment rate is clearly non-stationary, as indicated by high p-values in the Augmented Dickey-Fuller test at both lag lengths. It therefore needs to be transformed to achieve stationarity. The gift tax series initially signalled the presence of non-stationarity through the ADF test, which was later confirmed to be structural breaks with the Bai-Perron test. This suggests that the non-stationarity may arise from regime shifts rather than a stochastic trend. Finally, although the IVA time series appeared to be stationary, the Bai-Perron test detected breaks for which we might need to take appropriate measures.

4.4 Orders of Integration and Cointegration

Having established that the time series used in this study exhibit non-stationary behaviour (their statistical properties like mean, variance, and covariance change over time), applying traditional regression methods like OLS before appropriately treating the series risk producing spurious and misleading results. To determine whether multiple non-stationary series can be meaningfully modelled together, it is essential to examine their order of integration. A time series is said to be integrated of order zero, denoted I(0), if it is stationary in its level form. Instead, if a series is non-stationary in levels but becomes stationary after first differencing (subtracting each observation from its previous value), then it is considered to be integrated of order one, or I(1). Finally, if the series still remains non-stationary after first differencing but becomes stationary after taking the second difference, it is said to be integrated of order two, or I(2). This differencing process is essential in time series econometrics, as it enables the use of standard inferential procedures on data that initially violate the assumption of stationarity.

In the present thesis, differencing is a key step in identifying whether the series are integrated of order one, I(1), and thus eligible for consequent cointegration analyses. In econometrics, cointegration refers to a statistical property whereby two or more series have a common stochastic trend, meaning that while individual series may wander, they do so together in a way that maintains a stable long-run relationship. If this condition holds, we can introduce $Y_{t-1} - \theta X_{t-1}$ as an additional regressor called "error correction term", where θ is the cointegrating coefficient. Basically, if X and Y are cointegrated, computing the error correction term eliminates the common stochastic trend. In the context of this thesis, testing for cointegration allows us to determine whether variables like IVA or the gift tax co-move with income inequality over time in a meaningful economic sense, despite their individual non-stationarity.

For cointegration to be valid, all series involved must be integrated of the same order, typically I(1). If the series have different integration orders, such as a mix of I(1) and I(2) variables, standard cointegration methods like the Engle-Granger test or Johansen test cannot be applied. The testing strategy in this research thus began by computing the first difference of each time series using the *diff()* function in RStudio, followed by the Augmented Dickey-Fuller (ADF) test for unit roots to evaluate whether the differenced series were stationary. This procedure allows us to determine the order of integration of each variable and assess their eligibility for cointegration analysis.

The *ur.df()* function was applied to conduct the Augmented Dickey-Fuller (ADF) test for unit roots. This function allows to evaluate whether the variables exhibit non-stationary behaviour in their level or differenced forms. Thus, the test was initially applied to each variable in levels, which indicated non-stationarity across all series. Then, the same process was applied to each variable after first differencing. The results showed varying degrees of stationarity across the variables, but both income inequality and the value-added tax (IVA) series proved to be integrated of order one, I(1): first income inequality presented a test statistic of -3.057, exceeding the 1% critical value (-2.66) and indicating stationarity at the 1% level; second, the differenced value-added tax (IVA) variable produced a test statistic of -4.975, confirming it to be also I(1). By contrast, the test statistic for the differenced unemployment rate reported a value of -2.095, which lies between the 5% and 10% critical values. Although this value is some evidence of stationarity, it may not provide sufficiently strong evidence to confidently reject the null hypothesis at conventional significance levels (typically 5%). Finally, the differenced gift tax series yielded a test statistic of -1.239, which makes us fail to reject the unit root hypothesis at all conventional significance levels and suggests that the gift tax series remains non-stationary after first differencing. In fact, even after applying a second differencing, the gift tax series remained non-stationary. For this reason, the gift tax series was excluded from any further cointegration process.

Because the order of integration of the unemployment rate series can be a little ambiguous, and for purposes of clarity and transparency, the autocorrelation function (ACF) of the first-differenced series was first plotted, followed by theoretical reasoning for treating the series as integrated of order one, I(1).

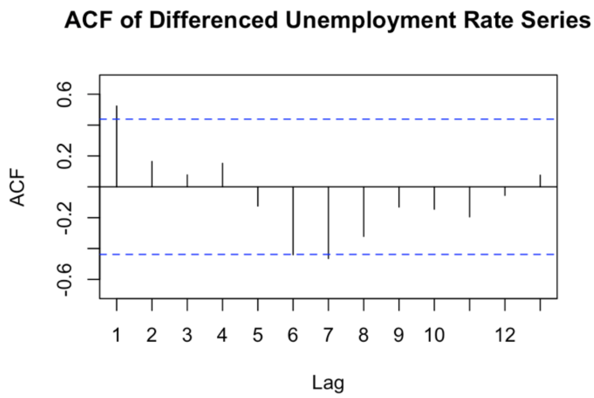

Table 12. (Source: Own elaboration)

Consistent with the previous results, Table 12 reveals a pattern consistent with weak stationarity: the autocorrelations drop off relatively quickly, and most lags fall within the 95% confidence bounds, suggesting the absence of a strong persistent trend and supporting the view that the series behaves as an I(1) process in practice. Moreover, there is theoretical precedent for treating this series as stationary in differences: in their empirical study of long-term U.S. time series, Nelson and Plosser (1982) sought to distinguish whether non-stationarity stemmed from stationary fluctuations around a deterministic trend or from stochastic processes with accumulating shocks. Their results favoured the latter, indicating that most macroeconomic series become stationary only after first differencing. Although Nelson and Plosser's (1982) analysis does not rule out the possibility of I(2) processes, their results still suggest that most macroeconomic time series, including the unemployment rate and real GDP, become stationary after first differencing. After having clarified the criteria employed to assess the order of integration of the variables, let us now proceed with the cointegration process. In this research, the Engle-Granger two-step procedure will be utilised.

The Engle-Granger test assumes that each individual series is integrated of order one, which is why it is necessary to compute the Augmented Dickey-Fuller (ADF) unit root test before proceeding with cointegration processes. Furthermore, it assumes that if the series are cointegrated, then their residuals should be stationary. The test thus involves two steps: estimating the long-run relationship among the differenced series through Ordinary Least Squares (OLS), and then checking whether the residuals from this regression are stationary using the ADF unit root test. In RStudio, the long-run equilibrium is estimated using the *lm()* function and incorporating the non-stationary series as *ineq_ts ~ iva_ts*. This computation means that the regression is estimating how much changes in income inequality are explained by the IVA series. The residuals resulting from this estimation represent the deviations from the estimated long-run path, which are to be tested for stationarity using the ADF unit root test.

	Engle-Granger Cointegration Regression
(Intercept)	0.233***
	(0.040)
	(<0.001)
iva_ts	0.019*
	(0.007)
	(0.011)
Num.Obs.	21
R2	0.295
R2 Adj.	0.258
F	7.957
RMSE	0.01

- $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Standard errors in parentheses.

- $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 13. Cointegration Regression for IVA and Income Inequality (Source: Own Elaboration)

As shown in Table 13, the estimated coefficient on IVA is positive, although moderately, with a value of approximately 0.019, and a p-value equal to 0.011. This indicates a modest, but still important long-run association between IVA and income inequality. statistical significance. These results provide preliminary evidence of a meaningful linear relationship between the two variables in levels. However, before confirming cointegration of the series, we must check for stationarity of residuals.

	ADF Test Regression
z.lag.1	-0.570*
	(0.233)
	(0.025)
z.diff.lag	0.225
	(0.255)
	(0.390)
Num.Obs.	19
R2	0.265
R2 Adj.	0.178
RMSE	0.01

- p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Standard errors in parentheses.

- p < 0.05, ** p < 0.01, *** p < 0.001.

ADF test statistic = -2.451

Critical values: 1% = -2.66, 5% = -1.95, 10% = -1.6

Table 14. ADF Test for Residuals of IVA and Income Inequality Cointegration Regression (Source: Own Elaboration)

As shown in Table 14, the resulting ADF unit root test statistic is reported to be -2.451. When compared against the critical value of -1.95 at the 5% significance level, the null hypothesis of non-stationary residuals is rejected. This confirms the presence of cointegration between IVA and income inequality. The result supports the existence of a stable long-run equilibrium relationship between the two variables, despite their individual non-stationarity. The same exact procedure was followed to test cointegration among the income inequality and unemployment rate time series.

	Engle-Granger Cointegration Regression
(Intercept)	0.321***
	(0.005)
	(<0.001)
unemp_ts	0.003***
	(0.001)
	(<0.001)
Num.Obs.	21
R2	0.563
R2 Adj.	0.540
F	24.521
RMSE	0.00

- p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Standard errors in parentheses.

- p < 0.05, ** p < 0.01, *** p < 0.001.

Table 15. Cointegration Regression for Unemployment Rate and Income Inequality (Source: Own Elaboration)

As Table 15 shows, the resulting estimated coefficient on the unemployment rate is very moderately positive. However, it is interesting to highlight the p-value (<0.001), suggesting a meaningful positive association between the two variables. Of course, to determine whether this relationship reflects a genuine equilibrium rather than a spurious correlation, the residuals from the regression were tested for stationarity.

	ADF Test Regression Test
z.lag.1	-0.752*
	(0.280)
	(0.016)
z.diff.lag	0.108
	(0.249)
	(0.670)
Num.Obs.	19
R2	0.341
R2 Adj.	0.264
RMSE	0.00

- p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Standard errors in parentheses.

ADF test statistic = -2.6851

Critical values: 1% = -2.66, 5% = -1.95, 10% = -1.6

Table 16. ADF Test for Residuals of Unemployment Rate and Income Inequality Cointegration Regression (Source: Own Elaboration)

As shown in Table 16, the test statistic is -2.685, which is more negative than the 5% critical value of -1.95. This result confirms the cointegration between income inequality and unemployment.

5. Error Correction Model (ECM) and Results

As was briefly explained in the previous chapter, a Vector Error Correction Model (VECM) is an advanced econometric framework which captures both the short-term dynamics and the long-term equilibrium relationships between multiple cointegrated series. When two or more variables share a common stochastic trend and are cointegrated, they tend to move together in the long run despite short-term deviations. Thus, a VECM augments a standard Vector Autoregression (VAR) model by incorporating an error correction term such as $Y_{t-1} - \theta X_{t-1}$, which reflects the previous period's deviation from the long-run equilibrium:

$$\Delta Y_t = \beta_{10} + \beta_{11}\Delta Y_{t-1} + \ldots + \beta_{1p}\Delta Y_{t-p} + \gamma_{11}\Delta X_{t-1} + \ldots + \gamma_{1p}\Delta X_{t-p} + \alpha_1(Y_{t-1} - \theta X_{t-1}) + \varepsilon_{1t},$$

$$\Delta X_t = \beta_{20} + \beta_{21}\Delta Y_{t-1} + \ldots + \beta_{2p}\Delta Y_{t-p} + \gamma_{21}\Delta X_{t-1} + \ldots + \gamma_{2p}\Delta X_{t-p} + \alpha_2(Y_{t-1} - \theta X_{t-1}) + \varepsilon_{2t}.$$

It is immediately noticeable that the VECM equation is the same as the VAR equation, only with the lagged differences of all the variables (expressed by the Δ) and the addition of the error correction term. Moreover, the coefficient α in the above equation indicates the speed of adjustment, that is, how quickly the system returns to equilibrium after a shock or disturbance. Thus, VECM can be a very useful and intuitive model, as it enables modelling of mutual interdependence and feedback effects among multiple variables. Unfortunately, because cointegration was tested using the two-step Engle-Granger method, a simpler Error Correction Model (ECM) will be employed instead. This is because the VECM is based on the Johansen methodology for cointegration, which was not employed in the present analysis. By contrast, the ECM is modelled on the Engle-Granger approach.

Despite it being simpler, the Error Correction Model (ECM) can still prove to be very useful. When cointegration occurs, it means that even though each variable may drift or trend over time, there exists a linear combination of them that converges to a stationary, long-term equilibrium. Moreover, as Enders (2014) explains, the time paths of cointegrated variables are influenced by any deviation from long-run equilibrium. This stationary combination implies the existence of a long-run equilibrium relationship. The error correction framework is thus designed to capture both short-run deviations and the long-run equilibrium relationship between such variables, as the short-term dynamics of the variables in the system are influenced by the deviation from such equilibrium (Enders, 2014).

Econometrically, the general structure of an Error Correction Model (ECM) can be expressed as follows:

$$\Delta Y_t = \beta_{10} + \beta_{11}\Delta X_{1,\,t-1} + \ldots + \beta_{1p}\Delta X_{1,\,t-p} + \gamma_{11}\,\Delta X_{2,\,t-1} + \ldots + \gamma_{1p}\,\Delta X_{2,\,t-p} + \alpha_1(Y_{t-1} - \theta_1 X_{1,\,t-1} - \theta_2 X_{2,\,t-1}) + \varepsilon_{1t}$$

Where ΔY_t is the first-differenced dependent variable (the income inequality series) integrated of order one; ΔX_1 and ΔX_2 are the first-differenced independent variables (the value-added tax (IVA) and unemployment rate series); and $Y_{t-1} - \theta_1 X_{1,\,t-1} - \theta_2 X_{2,\,t-1}$ is the error correction term. The coefficients $\beta_{11}, \ldots, \beta_{1p}$ and $\gamma_{11}, \ldots, \gamma_{1p}$ are the coefficients on the lagged first differences of X_1 and X_2, while the coefficient α_1 measures the speed of adjustment toward the long-run equilibrium.

For the sake of clarity, it is important to specify that in an Error Correction Model (ECM), the error correction term (ECT) expresses the long-run adjustment dynamics among cointegrated variables. The ECT is the lagged residual from the estimated long run cointegration equation (*ineq_ts ~ unemp_ts*, for instance), and it captures how much the system is out of equilibrium in the previous period. A negative and statistically significant coefficient on the ECT implies that the dependent variable adjusts in the current period to correct for deviations from the long-run equilibrium observed in the previous period. The first-differenced variables, instead, account for the short-run dynamics between variables. These terms reflect the immediate impact that changes in the explanatory variables have on short-term fluctuations in income inequality.

In RStudio, the model was estimated using the *lm()* function to calculate the linear regression between dependent and independent variable and catch the long-term equilibrium relationship among them. Finally, the Error Correction Model (ECM) was implemented in two versions: first, a bivariate model including the first difference of income inequality (*dy*), the first difference of IVA (*dx_iva*), and *ect_iva*; then a second bivariate model including *dy*, the first difference of unemployment (*dx_unemp*), and *ect_unemp*. This was done for the sake of consistency with the Engle-Granger methodology, which envisages the estimation of a single bivariate cointegration equation. Moreover, since both differencing and lagging operations introduce missing values at the beginning of the time series, the function *na.omit()* was applied to each dataset to ensure only complete and valid observations were included in the regressions.

	Error Correction Model
(Intercept)	0.000
	(0.001)
	(0.990)
dx_iva	0.008
	(0.006)
	(0.196)
ect	-0.389+
	(0.194)
	(0.061)
Num.Obs.	20
R2	0.223
R2 Adj.	0.131
AIC	-150.0
BIC	-146.0
Log.Lik.	79.016
F	2.433
RMSE	0.00

- p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table 17. Error Correction Model for IVA on Income Inequality (Source: Own Elaboration)

As Table 17 shows, the estimation on the differenced IVA variable is positive, although not statistically relevant due to a p-value of 0.194. Thus, short-term changes in IVA may not exercise an immediate effect on fluctuations in income inequality. Instead, the error correction term (ECT) tells a slightly different story: in fact, the ECT is negative and marginally significant at the 10% level (p-value = 0.061), suggesting that income inequality may exhibit a modest tendency to correct toward its long-run equilibrium level following a deviation, although weakly. Also, the negative sign of the ECT is consistent with econometric theory and implies that if income inequality rises above its long-run equilibrium level in one period, it will adjust downward in the subsequent period. Nevertheless, this first model's explanatory power is limited: the adjusted R-squared (which measures the goodness of fit of a statistical model) is relatively low with a value of 0.131. The overall F-statistic (which measures the overall significance of the model) presents a p-value of 0.118, meaning it cannot confirm joint statistical significance of the regressors.

	Error Correction Model
(Intercept)	-0.000
	(0.001)
	(0.905)
dx_unemp	0.001
	(0.001)
	(0.228)
ect	-0.642*
	(0.227)
	(0.011)
Num.Obs.	20
R2	0.348
R2 Adj.	0.271
AIC	-153.5
BIC	-149.6
Log.Lik.	80.773
F	4.532
RMSE	0.00

- p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table 18. Error Correction Model for Unemployment Rate on Income Inequality (Source: Own Elaboration)

Finally, Table 18 shows the second variation of the model, now including the first-differenced unemployment series and error correction term from the cointegration among unemployment rate and income inequality. The results reveal a statistically significant coefficient on the error correction term, which is presented with a value of -0.642, followed by a p-value of 0.012. This indicates that income inequality adjusts to deviations from its long-run equilibrium relationship with unemployment: in fact, the negative and significant ECT confirms the theory that when inequality is above its long-run path, there is a correction mechanism pulling it back towards equilibrium in the next period. In contrast, the short-run coefficient on the differenced unemployment variable (*dx_unemp*) is not statistically significant due to a p-value of 0.228, which tells us that changes in unemployment may not have an immediate short-term effect on income inequality. The goodness of fit indicator, the adjusted R-squared, is calculated at 0.271, with an F-statistic of 4.532 and overall p-value of 0.027. These results are moderate, but statistically significant. They reinforce the conclusion that long-run dynamics are driving the changes in income inequality in this case, rather than the short-run fluctuations.

In conclusion, the present results should be interpreted in light of the socioeconomic and institutional contexts that have characterised and plagued Italy for the past decades. For instance, Italy is characterized by a longstanding divide between the more prosperous North and the less developed South, where unemployment rates are systematically higher, economic activity is lower, and access to public services and infrastructure is often limited. These disparities mean that the aggregate national indicators employed in the present analysis as well, such as overall unemployment or income inequality, may obscure significant variations in how different regions experience and respond to macroeconomic shocks and policy interventions. These regional differences obviously complicate the interpretation of national-level econometric results. It is entirely plausible that the dynamics highlighted in the analysis may be stronger in certain regions than others.

An additional issue that is valid also on national scale is the rigidity of the Italian labour market. Research indicates that individuals in temporary positions often find themselves in a cycle of precarious employment, with limited opportunities for advancement (Bavaro & Tullio, 2024). Up to a certain degree, this situation is influenced also by rigid employment legislation, which, while righteously aiming to safeguard workers during recessions, it inadvertently backfires under better macroeconomic conditions (Boeri & Jimeno, 2016).

Taken together, these regional and institutional characteristics suggest that the findings of the present analysis, although moderate in their statistical significance, likely reflect deeper, historically rooted structural imbalances.

6. Limitations and Conclusion

The present analysis provided some insight into how fiscal policies and instruments and labour market factors in Italy explain changes in income inequality. This research was motivated by a long-standing interest in understanding the macroeconomic forces that shape a society and overall well-being in a specific socio-economic environment. For this purpose, applied statistics and econometrics provide reliable tools for empirical research, necessary to the obtainment of precise and quantitative data. Moreover, this thesis also represented a personal challenge to engage independently with more complex, quantitative methods, in support of personal academic and research interests.

Thus, this thesis thoroughly explained the process of variables selection through descriptive statistical analysis. Moreover, the inferential part of the thesis involved testing for unit roots to assess the stationarity properties of the time series, applying the Engle-Granger two-step method to identify cointegration relationships, and estimating multiple Error Correction Models (ECMs) to evaluate both short-term effects and long-run adjustments between income inequality and key macroeconomic indicators. In conclusion, the results suggest that neither IVA nor unemployment exert a statistically significant short-term effect on income inequality in the models tested. However, in one version of the model, the error correction term derived from the unemployment-based cointegration equation was significant, indicating the presence of a long-run equilibrium adjustment mechanism.

From a public policy perspective, this implies that policies focused on short-term effects on income inequality may have limited effects or be entirely non-effective. However, the presence of a long-run relationship between the unemployment rate and income inequality suggests that policies should focus on solving the structural unemployment issues in the country. Policymakers should then prioritize tailoring long-term, impactful policies instead of aiming for rapid redistribution effects. However, like all empirical analyses, this study also faces several constraints, the most important one being the lack of familiarity with advanced econometric tools. As specified earlier in the paragraph, this analysis represented a personal challenge. Naturally, this implied coming to terms with several technical limits.

For instance, the Engle-Granger and ECM method is a valuable, but limited tool for addressing relationships between cointegrated time series. As Enders (2014) notes, “the simplicity of simulated data is rarely encountered in applied econometrics,” which reflects the difficulties encountered by this methodology. Undoubtedly, it provides a useful starting point. Just that its assumptions and structure may not fully capture the complexity and interdependence of macroeconomic variables such as inequality, taxation, and unemployment. Yet, the Johansen system-based approach, although offering a more detailed analysis of multivariate cointegration, represents a level of statistical complexity and analysis that stray beyond the feasibility of this thesis. Still, the Engle-Granger method is a respected technique due to conceptual clarity and accessible interpretation. This decision reflects a deliberate trade-off between statistical complexity and accessibility, allowing for an in-depth exploration of both long-run equilibrium relationships and short-run adjustments within a more accessible econometric framework.

7. Bibliography

Assous, M., & Carret, V. (2022). Moving dynamics beyond business cycles: Jan Tinbergen's first macrodynamic model (1934–1936) . *The European Journal of the History of Economic Thought* , 117-131.

Basso, F. (2016, May 6). *L'Europa e la crisi del debito sovrano, ecco le tappe dalla Grecia all'Italia*. Retrieved from Corriere: https://www.corriere.it/economia/cards/europa-crisi-debito-sovrano-ecco-tappe-grecia-all-italia/novembre-2010-sistema-bancario-irlandese-crisi.shtml

Baum, C. F. (2013). *EC823: Applied econometrics: Lecture 10—Cointegration and error correction models.* Retrieved from Boston College: http://fmwww.bc.edu/EC-C/S2013/823/EC823.S2013.nn10.slides.pdf

Bavaro, M., & Tullio, F. (2024). A cycle or a tunnel? A study on unemployment and low-pay dynamics in Italy. *Labour Economics, 90*.

Berg, A. G., & Ostry, J. D. (2011). *Inequality and Unsustainable Growth: Two Sides of the Same Coin?* International Monetary Fund.

Blanchard, O., & Perotti, R. (2002). An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *the Quarterly Journal of economics, 117 (4)*, 1329–1368.

Boeri, T., & Jimeno, J. F. (2016). Learning from the Great Divergence in unemployment in Europe during the crisis. *Labour Economics, 41*, 32-46.

Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time Series Analysis: Forecasting and Control.* Hoboken, New Jersey: John Wiley & Sons Inc.

Cicchitelli, P. D. (2021). *Statistics: Principles and Methods.* Pearson.

Enders, W. (2014). *Applied Econometric Time Series (4th ed.).* Wiley.

Fisher, R. A. (1924). On a distribution yielding the error functions of several well known statistics. *Proceedings International Mathematical Congress, Toronto*, 2, pp. 805-813.

Galton, F. (1886). Regression towards mediocrity in hereditary stature. *Journal of the Anthropological Institute of Great Britain and Ireland 15*, 246–263.

Haavelmo, T. (1944). The Probability Approach in Econometrics. *Econometrica*, pp. iii-vi, 1-115.

Klein, L. R. (1950). *Economic Fluctuations in the United States, 1921-1941.* New York: J. Wiley and Sons.

Lucas, R. E. (1976). Econometric policy evaluation: A critique . *Carnegie-Rochester Conference Series on Public Policy, 1*, 19-46.

Ministri, P. d. (n.d.). *4. Occupazione*. Retrieved from Dipartimento per la programmazione e il coordinamento della politica economica: https://www.programmazioneeconomica.gov.it/it/focus/andamenti-di-lungo-periodo-dell-economia-italiana/grafici/4-occupazione/

Moore, H. L. (1911). *Laws of Wages: An Essay in Statistical Economics.* New York: MacMillan.

Moran, P. A., & Whittle, P. (1951). Hypothesis testing in Time Series Analysis. *Journal of the Royal Statistical Society. Series A (General), Vol. 114, No. 4*, 559.

n.d. (2025, May 6). *8.3 Autoregressive models | Forecasting: Principles and Practice (2nd ed)*. Retrieved from https://otexts.com/fpp2/AR.html

Nelson, C., & Plosser, C. (1982). Trends and random walks in macroeconmic time series . *Journal of Monetary Economics 10*, 139-162.

Schilirò, D. (2014). *Crescita e debito pubblico. L'economia italiana e la crisi.* Munich Personal RePEc Archive.

Sims, C. A. (1980). Macroeconomics and Reality.

Stock, J. &. (2020). *Introduction to Econometrics (4th ed.).* Pearson.

Walker, G. T. (1931, June). On periodicity in series of related terms . *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, pp. Volume 131, Issue 818.

Yule, G. U. (1927, January). On a method of investigating periodicities disturbed series, with special reference to Wolfer's sunspot numbers. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, pp. Volume 226, Issue 636-646.