The Economic and Social Review, Vol. 50, No. 3, Autumn 2019, pp. 429-458

Examining the Volatility of Ireland's Tax Base in the Paradigm of Modern Portfolio Theory

Keith Fitzgerald*
Parliamentary Budget Office, Dublin

Jacopo Bedogni
Parliamentary Budget Office, Dublin

Abstract: This paper extends a theoretical and empirical framework commonly applied in international finance to present an alternative paradigm within which issues of tax revenue volatility in Ireland can be studied. We establish a trade-off between revenue growth and volatility, typically associated with financial asset returns. We observe strong but time-varying cointegration among tax revenue streams. Statistical tests of mean-variance spanning suggest diversification benefits from holding Income Tax and Excise Duty. We establish the minimum variance tax portfolio, and also find that, from a mean-variance optimality perspective, the portfolio consisting of 2017 tax shares is sub-optimal. Practical policy implications are discussed.

I INTRODUCTION

Prudential fiscal planning necessitates that long-term public spending be based on stable and less volatile sources of revenue. If a programme of expenditure is funded using windfall tax receipts, additional revenue (or spending cuts to other areas) will be required to continue to fund this programme should these windfall receipts fail to materialise in future years. In addition, revenue volatility complicates fiscal planning, as revenue forecasting is likely to be more challenging for more volatile taxes.

Acknowledgements: The authors would like to thank Martha O'Hagan Luff, Kieran McQuinn, Members of the Parliamentary Budget Office's External Advisory Group, and two anonymous reviewers for constructive feedback and advice. The authors are solely responsible for the content and the views expressed.

* Keith.Fitzgerald@Oireachtas.ie or Jacopo.Bedogni@Oireachtas.ie.

For ease of fiscal planning, and to ensure that public spending is linked to sustainable sources of revenue, a state might be interested in rebalancing the tax base and setting the share of each tax category at levels that minimise total tax revenue volatility. In other words, this would be the volatility-minimising tax base composition. Furthermore, an understanding of the "risk and return trade-off" for each of the state's main taxes more generally is valuable information for fiscal policymakers to have.

Of course, the choice of tax mix does not only concern the goal of minimising volatility, but also includes wider issues such as equity and efficiency. Consideration should be given to how taxes impact on growth (some taxes are more distortionary than others), or influence the behaviour of tax-payers (whether intentional or otherwise). In effect, the optimal tax shares from the perspective of portfolio theory might differ substantially from the actual or desired tax shares when these wider issues are taken into account in designing the tax system (including feasibility constraints and social preferences). However, these issues have been well researched already (for a recent literature review of the evidence see Johansson, 2016) and are beyond the scope of our analysis. Instead, our paper aims to make a novel contribution by providing a theoretical discussion of how the State's portfolio of taxes might be structured if guided solely by the principles of risk and return.

We apply Modern Portfolio Theory to present an alternative paradigm within which issues of volatility and growth in respect of State tax revenue can be studied. We focus specifically on the seven largest categories of tax that comprise Exchequer revenue (Income Tax, VAT, Corporation Tax, Excise Duty, Stamp Duty, Capital Taxes and Customs). We compare the share of total tax revenue for each of these taxes to the financial asset shares in an investor's portfolio. In our analysis, we replace financial asset returns with the "returns" of each tax category, that is; the revenue growth for each tax. Similarly, the volatility or risk of each tax category is assessed as the standard deviation of these "returns" (a standard measure of volatility or risk in financial economics).

Generally the volatility of an investment portfolio is a function of the volatility of the underlying financial assets, their share in the investment portfolio and the covariance between these assets. Likewise, we model the overall volatility of total tax revenue as a function of:

1. the share of each tax category in the State's tax portfolio;
2. the volatility of each individual tax category; and
3. the covariance and correlation among these tax categories.

Intuitively, the higher the volatility of the individual revenue streams comprising the State's tax portfolio, and the more positive the correlation among these revenue streams (i.e. the extent to which they co-move); the greater total tax revenue volatility will be.

Our results have practical implications for taxation policy in Ireland. We argue generally, that there is greater scope for the consideration of the impact on overall revenue volatility when designing policies that could materially alter the composition of Irish Exchequer revenue. We establish that there is a risk-return trade-off for the primary categories of tax, and that overall tax portfolio volatility is lower when the portfolio is concentrated around fewer, but less volatile taxes. From a diversification perspective, we observe strong (time-varying but consistent) cointegration, implying limited benefits from diversifying across the main tax categories.

Furthermore, we establish the shares of each tax that minimise overall tax portfolio volatility. We propose that these shares could act as benchmark shares. These shares could guide allocations to a "Rainy Day Fund" (RDF) – based on the extent to which the share of a volatile tax in a given year exceeds the estimated benchmark share – or could be used to steer incremental policy change towards minimising the overall volatility of Exchequer revenue. We emphasise the importance of counter-cyclical (or Rainy Day) funds in mitigating the risks of over-reliance on volatile revenue streams, and in ensuring funding for expenditure commitments when volatile revenues fail to materialise in future years.

This paper is divided into six main sections. Section II outlines the context to this research. In Section III we give an overview of the relevant literature, while in Section IV we provide a preliminary analysis of the underlying data. Section V illustrates the statistical methods used in this paper and Section VI details our empirical analysis. Finally, in Section VII we conclude and discuss relevant policy implications, and outline avenues of potential future work.

II RESEARCH CONTEXT

The risks associated with the volatility of Ireland's tax base (i.e. large yearly swings in government revenue), particularly in relation to Corporation Tax receipts, have come to the fore of debate on fiscal policy in Ireland. Having a tax base that consists of highly volatile revenue streams can potentially undermine the stability and sustainability of the public finances. Expenditure commitments of a permanent nature (e.g. public sector pay and pensions) should be linked to stable and less volatile revenue sources. If a new spending programme is funded using windfall tax receipts, additional revenue or spending cuts to other areas will be required to continue to fund this programme should these windfall receipts fail to materialise in future years. Furthermore, revenue volatility complicates fiscal planning, given that forecasting is likely to be more difficult for more volatile taxes. This paper investigates these issues and provides an analysis of the volatility of Ireland's tax base within the paradigm of Modern Portfolio Theory.

In our view, Modern Portfolio Theory can be generalised to an analysis of a state's tax revenue. In Modern Portfolio Theory, which is drawn from financial economics, an investor has a portfolio made up of different financial assets, each of which has a different risk and expected return profile. The relationship between risk and return in relation to financial assets is intuitive; riskier investments will need to compensate would-be investors for taking on additional risk relative to alternative safer investment options. More generally, if a low-risk high-return asset did exist, rational investors would flock to that asset, driving up the price and reducing the rate of return to a level that the market deems more appropriate to that level of risk. The reverse would hold for a high-risk low-return asset.

If we consider a two-asset portfolio, the overall portfolio variance is calculated as follows:

$$V_p = (w_a^2\sigma_a^2) + (w_b^2\sigma_b^2) + 2(w_a\sigma_a w_b\sigma_b cov(a,b)) \tag{1}$$

where w_a is the portfolio share of the first asset a; w_b is the portfolio share of the second asset b; σ_a is the standard deviation of asset a; σ_b is the standard deviation of asset b; and $cov(a,b)$ is the covariance between the two assets. The volatility of the portfolio is calculated as the standard deviation of the portfolio (the square root of the variance of the portfolio).

Generally, an investor would pursue two main investment strategies:

1. choosing the asset mix that minimises portfolio volatility for a given return; or
2. choosing the asset mix that maximises the portfolio return for a given level of volatility.

Those portfolios that represent the maximum return for a given level of risk (or the minimum level of risk for a given level of return) are said to lie on the mean-variance efficient frontier. These portfolios are mean-variance efficient, and there is no other portfolio that is more appealing within the investor's opportunity set.[1]

Our empirical analysis is three-fold:

- We first examine the degree of correlation and cointegration (standard measures of the level of integration between financial market series) among the main tax categories comprising the State's tax portfolio. Theoretically, the more integrated the revenue streams are, the smaller the diversification benefits from holding a portfolio consisting of these revenue streams.

[1] There is no obvious revenue source that could serve as a counterpart to a risk-free asset in the context of Exchequer revenue. Therefore, we consider a portfolio consisting of risky assets only. The efficient frontier consists of the portfolios that optimally combine the set of risky assets, in terms of risk and return (see Markowitz, 1952).

- We then construct the mean-variance efficient frontier for a portfolio of all taxes (which begins with the global minimum variance portfolio). In practice, this involves minimising the portfolio variance equation in (1) subject to the constraint that the share of taxes in the portfolio be non-negative and sum to one, for various levels of return. We subsequently construct the efficient frontier for a portfolio of all taxes minus one, and compare this to the frontier consisting of the complete set of taxes. This enables us to visually examine the impact (if any) on the frontier from adding the missing category of tax, and we repeat this for each tax.
- Finally, we use statistical tests of mean variance spanning to assess if the change in the efficient frontier from adding the missing tax is statistically significant. This methodology allows us to examine the potential diversification benefits that each individual tax category brings to the State's portfolio of taxes, in terms of risk and return.

III LITERATURE

This section examines the relevant literature, and outlines how our paper complements and expands on this literature.

Relatively few papers have studied the volatility of tax revenue within the paradigm of Modern Portfolio Theory. Existing studies tend to focus on the volatility of state taxes in the US, given the Federal governance system of the US and the extent to which fiscal policies are determined at state level. However, these studies do contain useful learnings that can be applied to the Irish context.

Crain (2003) is among the earliest papers to apply financial market theory to an analysis of tax revenue volatility at state level in the US, examining the trade-off between the growth rate of tax revenue and the volatility or riskiness of that revenue. This is commonly referred to as the risk-return trade-off in the context of financial markets. Our paper establishes the existence of a trade-off in the Irish context also, and further compares the risk and return profile of the primary taxes to investments in US Treasury Bills (T-Bills), the S&P 500, and an index of emerging market equities.

Matthews (2005) considers the funding of permanent expenditure using a volatile revenue source. The paper calculates the coefficient of variation (the ratio of the standard deviation and revenue growth) and prediction errors for property, sales and income taxes in the US state of Georgia. In particular, Matthews (2005) argues for the diversification of funding for state-wide education across multiple revenue sources, given the volatility inherent in a single tax. The paper further recommends building budgetary reserves in times of growth to preserve funding during economic downturns. We similarly examine the coefficient of variation as part of our analysis of the seven primary Irish taxes, and we also argue for placing more windfall revenue from volatile revenue sources into a counter-cyclical fiscal fund.

Garrett (2006) compares the actual revenue share of a tax to its variance minimising share, and examines how well a state's portfolio is constructed to minimise the variance in total tax revenue for the state. He argues that future work could examine the stability of variance minimising shares over different points of the business cycle, and proposes extending the model to allow for the analysis of multiple taxes rather than considering a specific tax against a combination of all "other" tax revenue. Our paper builds on this analysis by examining the variance minimising shares for multiple taxes (i.e. for all seven taxes in our sample), and across multiple time periods.

Cornia and Nelson (2010) conduct a similar assessment across US states, with each state representing a different portfolio of taxes. They construct the efficient frontier and examine how the actual composition of tax revenue across states compares with respect to the efficient frontier. From this, they argue that consideration should be given to the volatility and growth of a tax portfolio when designing tax policy changes that will impact on the composition of overall state revenue. We also examine how Ireland's tax portfolio compares to the efficient frontier, and further examine how the efficient frontier is impacted from the addition of each individual tax category, with both visual and statistical tests of mean-variance spanning.

McQuinn and Roche (2016) are the first to apply mean-variance analysis in the Irish (and indeed European) context. They examine Irish Exchequer receipts from 1984 to 2015 and estimate the efficient frontier based on *ex ante* tax forecasts and *ex post* outturns. Their analysis shows that an improvement of tax revenue growth relative to revenue volatility was possible if a larger proportion of tax revenue was derived from direct taxes (e.g. Income Tax) rather than indirect taxes (e.g. VAT). While we similarly analyse the efficient frontiers, we carry out statistical tests of mean-variance spanning and measure the degree of cointegration among tax revenue streams in a dynamic framework. We also estimate the variance minimising shares for a more comprehensive portfolio of taxes including Capital Taxes and Stamp Duty. This is important, as each of these tax categories has a unique risk-return profile, and the behaviour of these taxes throughout the fiscal crisis in particular necessitates that they be studied in separation.

There is a relatively larger body of literature in the wider area of tax elasticity and revenue buoyancy. These studies model revenue volatility as the elasticity of tax revenue with respect to changes in some macroeconomic variable, often GDP, and seek to determine the extent to which tax revenue moves with the business cycle. Pro-cyclical taxes tend to outperform expectations and experience strong outturns in times of growth, while the reverse is true during an economic downturn. As an example of this work in the Irish context, see Deli *et al.* (2018) and Acheson *et al.* (2017). As an alternative, our analysis focuses on the standard deviation of the revenue growth series as a measure of revenue volatility (as is standard practice in Modern Portfolio Theory). Additional work has examined the implications of tax revenue volatility for revenue forecasting in Ireland (for example, Hannon, 2014, analyses the source of forecast errors in an error decomposition framework).

IV PRELIMINARY ANALYSIS

4.1 Data

We use annual tax revenue data taken from the Databank of the Department of Public Expenditure and Reform, which includes outturn data for the following tax categories:

- Income Tax;
- Value Added Tax;
- Excise Duty;
- Corporation Tax;
- Stamp Duty;
- Customs; and
- Capital Taxes (Capital Gains Tax and Capital Acquisitions Tax).

The use of annual tax data means that we somewhat mitigate the issues relating to the seasonality of tax receipts (i.e. filing and consumption patterns). We focus on Exchequer revenue sourced from tax receipts that are used to fund the vast majority of General Government expenditure. Most importantly, these revenues are not ring-fenced for a particular purpose. Specifically, we exclude from our calculations the Training and Employment Levy, PRSI receipts, the Local Property Tax, and Motor Vehicle Duties, as these revenues are generally ring-fenced to fund specific expenditure. For consistency purposes, Capital Gains Tax and Capital Acquisitions Tax are aggregated and defined as Capital Taxes (these taxes have only been reported separately in outturn data since 2000). We use this revenue outturn data to calculate revenue growth rates for each tax and the standard deviation of these growth rates.

4.2 The Revenue "Risk and Return Trade-Off"

Table 1 shows the mean return and standard deviation, as well as the ratio of the two, for the seven taxes in our sample from 1985-2017. We observe that Capital Taxes and Corporation Tax tend to have the highest average annual growth rates, and, alongside Stamp Duty, are among the most volatile in terms of standard deviation.

Among the largest taxes (Income Tax, VAT, Excise Duty and Corporation Tax) Corporation Tax has the highest standard deviation (17.3 per cent), meaning that the annual percentage changes in Corporation Tax would tend on average to fluctuate significatly from their mean value (17.3 percentage points above or below the average annual change). Corporation Tax also exhibits the highest average annual growth rate at 12.2 per cent, while Excise Duty grew on average by 4.2 per cent each year over the period. Excise Duty is the tax category with the lowest standard deviation (5.6 per cent). Income Tax, VAT, and Excise Duty generally perform well in terms of their mean return or growth, relative to their volatility (0.96, 0.84 and 0.76 respectively).

The data in Table 1 suggest the existence of a trade-off in terms of the rate of tax revenue growth (or return) and volatility, more typically associated with assets traded in financial markets. This trade-off is clearer from Figure 1, which for each tax category, plots the standard deviation against the average growth rate. In Figure 1, the risk-return trade-off is assessed relative to investments in US T-Bills (low volatility and low return investments), the S&P 500 (medium volatility and medium return investments), and an index of emerging market equities (high volatility and high return investments). Generally (with the exception of Corporation Tax), revenue streams having high volatility and large returns, such as Capital Taxes and Stamp Duty, are pro-cyclical, being based on transactions involving activities that are subject to "boom and bust" cycles.

Table 1: Mean Revenue Growth and Standard Deviation, 1985-2017

Tax Category	*Mean Growth Rate (Return)*	*Standard Deviation (Risk or Volatility)*	*Return-Risk Ratio*
Income Tax	6.61%	6.90%	0.96
VAT	6.68%	7.91%	0.84
Excise Duty	4.24%	5.55%	0.76
Corporation Tax	12.24%	17.27%	0.71
Stamp Duty	9.42%	23.07%	0.41
Customs	4.25%	14.96%	0.28
Capital Taxes	16.38%	33.91%	0.48

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank.
Notes: Capital Taxes comprise Capital Acquisitions Tax and Capital Gains Tax (reported separately in outturn data since 2000).

4.3 Ireland's Tax Revenue Structure Over Time

As outlined previously, portfolio variance is a function of the volatility of the underlying assets, their share in the portfolio and the covariance between these assets.

Figure 2 provides an overview of the trends in tax revenue shares from 1984-2017. Historically, Income Tax, VAT and Excise Duty have accounted for the largest proportion of total tax revenue (with a combined annual average from 1984-2017

Figure 1: Risk-Return Trade-Off among the Main Tax Categories, 1985-2017

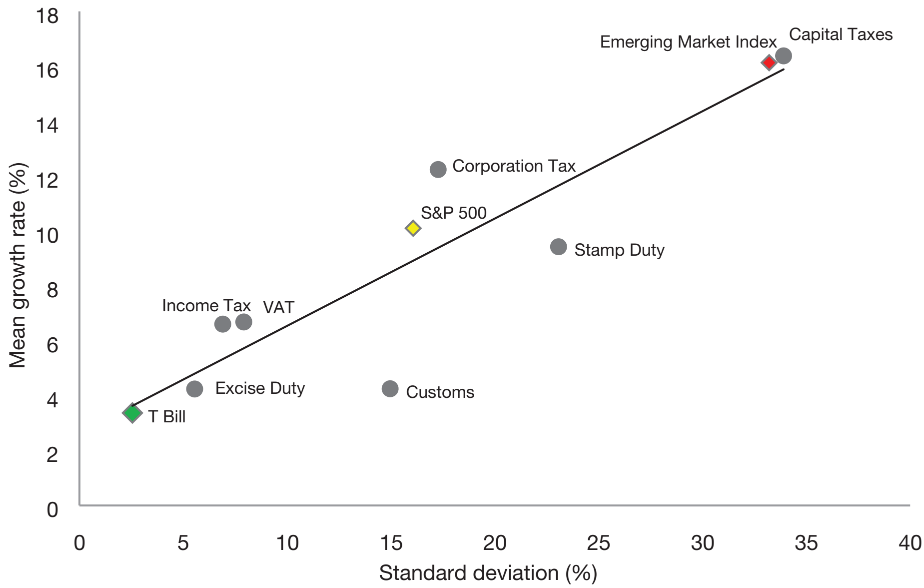

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank. T Bill rate taken from US Treasury, S&P 500 and Emerging Market Index taken from MSCI.

Figure 2: Trends in Tax Shares Over Time, 1984-2017

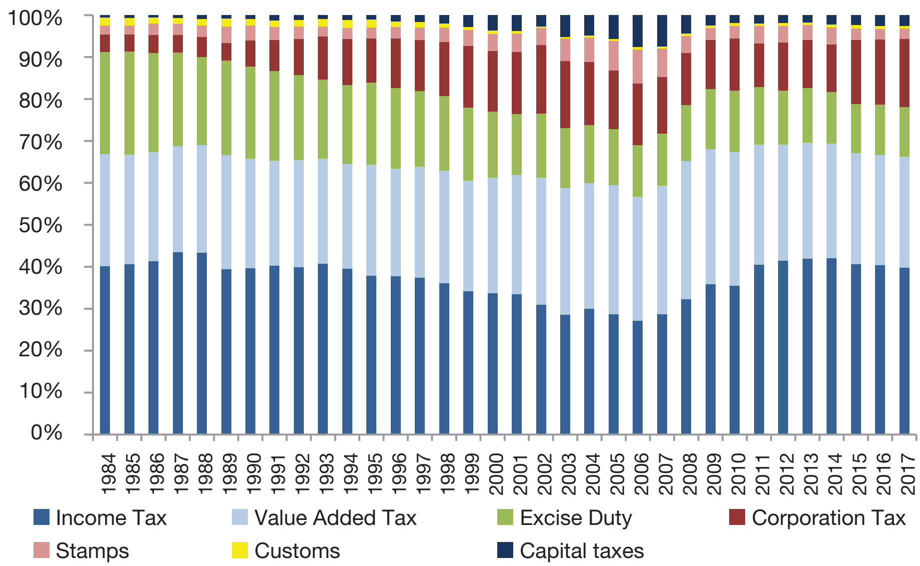

Source: Authors based on Department of Public Expenditure and Reform – Databank.
Note: The shares above are calculated in respect of total revenue.

of 80.2 per cent). However, the share of Corporation Tax (16.2 per cent in 2017) has recently surpassed the share of Excise Duty (11.7 per cent in 2017), becoming the third most important revenue stream for the State. In recent years, Corporation Tax receipts have surged, increasing by 95 per cent from 2012 to 2017, with a 49 per cent increase in 2015 alone. This may have been influenced by the decision taken by foreign-owned multinationals to relocate their assets and activities to Ireland. Corporation Tax revenue accounted for 16.2 per cent of total tax revenue in 2017 (or €8.2 billion), which is well above the long term-average over 1984-2016 of 10.8 per cent.

The trend in Corporation Tax receipts emphasises the inherent volatility of this tax category, due to the responsiveness of receipts to wider economic conditions and to changes in tax regimes, both internationally and domestically. Corporation Tax receipts are also very concentrated, with data from the Revenue Commissioners (2018) highlighting that in 2017:

- the top ten tax-payers accounted for 39 per cent of receipts; and
- foreign-owned multinationals accounted for 80 per cent of receipts.

These findings highlight the extent to which Corporation Tax receipts are exposed to firm and sector-specific shocks.

The pre- and post-crisis period saw significant changes to Income Tax, which impacted on its relative share. Income Tax is the largest revenue stream of the Exchequer. This is true both historically and at present, with taxes on income (including USC) accounting for 39.4 per cent (or €20 billion) of Exchequer revenue in 2017 (relative to an average of 36.4 per cent over 1984-2016). The Income Tax share fell to a low of 27.2 per cent in 2006. This is partially explained by the pre-crisis implementation of budget measures that narrowed the Income Tax base, and by the dramatic rise in VAT revenue during this time (Income Tax receipts grew by 36 per cent from 2000-2006, while VAT receipts grew by 80 per cent over the same period). In the aftermath of the fiscal crisis of 2008, in an attempt to restore the public finances, Income Taxes were increased and several tax expenditures were reduced. As a result, the Income Tax share peaked at around 42 per cent in 2013 and 2014. In recent years, Income Tax receipts have grown strongly (increasing by 27 per cent since 2013), notwithstanding the implementation of measures which have contributed to the narrowing of the base. In fact, Income Tax shows a significant degree of concentration with the top 13 per cent of total Income Tax units (those with gross income above €70,000), accounting for 64 per cent of total tax paid (including USC) in 2016 (based on the Revenue Commissioners' Income Tax and Corporation Tax distribution statistics).

Similarly to Income Tax, VAT receipts have increased by 28.7 per cent since 2013. VAT is the second largest revenue stream accounting for 26.2 per cent of total tax receipts (or €13.3 billion) in 2017, below the long-term average over 1984-2016 of 27.2 per cent. However, the VAT share in total tax revenue has decreased in recent years, due to the increasing share of Corporation Tax.

During the boom period, the share of Stamp Duty in total tax revenue doubled from 4.1 per cent in 2000 to 8.2 per cent in 2006, before falling to 2.8 per cent in 2009 (with a share of 2.4 per cent in 2017). Similarly, for Capital Taxes, the share increased from 3.7 per cent in 2000 to 7.6 per cent in 2006, before falling to 2.4 per cent in 2009 (with a share of 2.5 per cent in 2017).

4.4 Diversification of the Revenue Base

Portfolio diversification has the potential to reduce portfolio volatility, assuming diversification across holdings that are not significantly positively correlated. Figure 3 shows the degree of concentration or diversification of tax revenue over time, alongside changes in the standard deviation of the State's tax portfolio. In particular, we are interested in assessing whether Ireland has diversified its tax portfolio over time, and how this compares with portfolio volatility.

Figure 3: Tax Portfolio Diversification and Volatility, 1984-2017

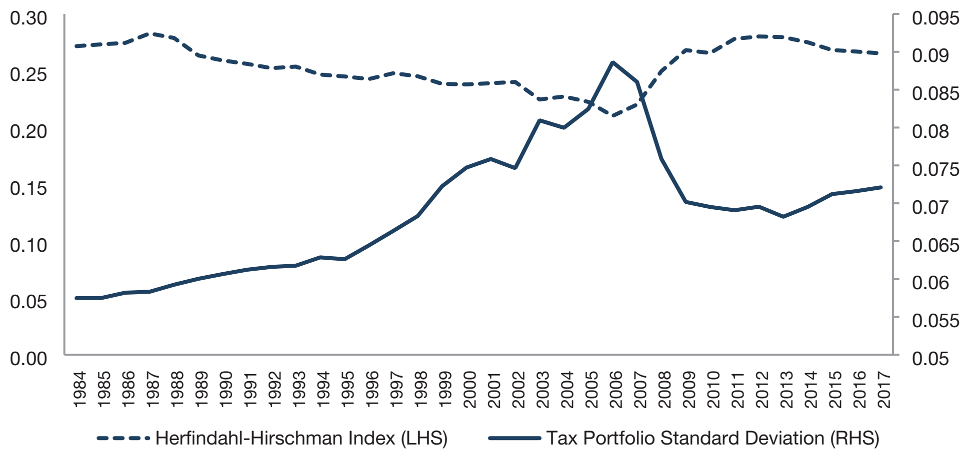

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank.

Notes: The HHI ranges from 0 to 1, with 1 being the case that all tax revenue is coming from a single source. As of end-2017, the tax base is largely concentrated around Income Tax, Excise Duty, Valued Added Tax and Corporation Tax. These taxes make up 94 per cent of the revenue aggregate across the taxes included in our sample.

We calculate the level of concentration in the tax portfolio using the Herfindahl–Hirschman Index (HHI), which is an index typically employed in competition economics to measure the level of market power in an industry. While this approach is generally applied to market shares of firms, we extend it to an analysis of Ireland's tax shares. The index ranges between 0 and 1, with 1 being the case that all tax revenue is coming from a single source. Thus, the higher the index, the lower the degree of diversification of the tax base (i.e. the base is highly concentrated in a small number of taxes). The mathematical representation of the HHI index is as follows:

$$HHI = \sum_{i=1}^{N} w_i^2 \tag{2}$$

where w_i is the relative tax share of a given tax category i, and N is the number of tax categories. Generally, financial market theory would suggest that more diversification is better in terms of risk reduction. However, this must take account of the co-movements among tax revenues and the underlying volatility of each revenue stream. The dotted line in Figure 3 highlights that historically Ireland's tax base has been highly concentrated in a small number of taxes. From 1989 the tax base became more diversified, with the highest level of diversification reached in 2006. Since then, the trend has reversed and the tax base has become more concentrated.

The solid line in Figure 3 highlights an interesting point: the increase in the diversification of tax revenue was associated with an overall increase in portfolio volatility or risk (we explore this in more detail in Section 5). The inverse relationship between the two series (with a correlation of -71 per cent) was particularly evident during the boom period. This can be explained by the tax base shifting towards more volatile revenue sources (i.e. an increasing share of Capital Taxes and Stamp Duty). Addison-Smyth and McQuinn (2010; 2016) find that the period from 2002 to 2009 is characterised by substantial windfall gains in Stamp Duty and VAT, above levels warranted by underlying economic fundamentals. In this way, while tax revenue was seemingly more diversified during this time, this spike in the share of certain taxes was temporary and unsustainable. With the 2008 economic and financial crisis, the inevitable collapse of these revenues, and the fiscal consolidation measures that followed, greater emphasis was placed on relatively more stable sources of revenue, such as Income Tax.

V STATISTICAL METHODS

5.1 Correlation and Cointegration of Tax Revenue

5.1.1 Correlation

We examine the degree of linear association among the revenue streams of the seven tax categories in a bivariate framework using Pearson's correlation coefficient:

$$r = \frac{n(\Sigma xy) - (\Sigma x)(\Sigma y)}{\sqrt{[n\Sigma x^2 - (\Sigma x)^2][n\Sigma y^2 - (\Sigma y)^2]}} \tag{3}$$

The correlation coefficient is bound above by 1, and below by –1 (with 1 indicating a perfect positive correlation and –1 indicating a perfect negative correlation). When a pair of variables moves together in the same direction, there is a positive correlation between these variables. From an investment perspective, a portfolio consisting of underlying assets that exhibit strong positive co-movement (correlation) is undesirable, as an external shock will impact on the performance of all of the assets in the portfolio in a similar way. The same is true of a portfolio consisting of the State's tax revenues. Portfolio diversification generally involves allocating shares across entities that do not exhibit strong positive correlation with one another.

5.1.2 Cointegration

Following an analysis using Pearson's correlation coefficient, we model the long-term linkages among the seven tax categories as cointegrating relations using the Johansen-Juselius test (1990). This is a commonly used method in assessing the level of integration within a system or between a pair of variables, and has been employed extensively in studies of financial market integration (for example, see Quayes and Jamal, 2016; Neaime, 2016; and Babalos *et al.,* 2016). The Johansen-Juselius test is predicated on an assumption that the series contain a unit root, which we have pre-determined using an Augmented Dickey-Fuller test.

Essentially, the Johansen-Juselius cointegration test involves calculating the test statistics for a vector auto-regressive (VAR) model in which each of our seven tax categories act as dependent variables. The resulting test statistics are used to identify the number of cointegrating vectors in the system defined by the VAR. The Johansen-Juselius test produces two test statistics. While the trace statistic examines the null hypothesis of r cointegrating vectors against the alternative hypothesis of n cointegrating vectors, the maximum eigenvalue statistic considers the null hypothesis of r cointegrating vectors against the alternative of $r + 1$ cointegrating vectors. Johansen and Juselius (1990) provide asymptotic critical values.

We examine the level of cointegration in the complete set of taxes in a time-varying recursive framework, allowing for intertemporal changes in the extent to which tax revenue streams are integrated (following Hansen and Johansen, 1993; Lucey and Voronkova, 2005). This recursive approach involves estimation over an initial period of size t (in our case this is the 16 years from 1985 to 2000) before expanding this period incrementally (by a single year) and re-estimating, and continuing until the end of the sample. We then plot each of the corresponding trace statistics, scaled to the 90 per cent critical value, which allows us to observe the dynamics of cointegration among the tax revenue streams over time.

5.2 Tests of Mean-Variance Spanning

Investment theory has long advanced the view that investors should seek to diversify their holdings to reduce their risk exposure (see Markowitz, 1952). Modern Portfolio Theory is underpinned by the idea that the investment decision should be based not only on the risk-return profile of an individual asset, but on the relationship between that asset and the various other assets that the investor wants to hold.

Tests of mean-variance spanning (MVS) are used extensively in portfolio analysis to examine the implications of expanding a portfolio of assets, to include additional assets, in terms of the risk-return trade-off. Spanning tests were introduced by Huberman and Kandel (1987) and further developed by Kan and Zhou (2008), and have dominated the international finance literature since (see Berrill and Kearney, 2010; Eiling *et al.,* 2012; O'Hagan-Luff and Berrill, 2015).

MVS analysis considers a set of K "benchmark" and N "test" assets (typically financial assets), and investigates if, conditional on the set of K benchmark assets, the addition of N test assets can shift the mean-variance efficient frontier. In other words, MVS involves examining if the efficient frontier of the benchmark set coincides with the efficient frontier of the extended set (K+N). This amounts to examining if the set of benchmark assets yields the same diversification benefits as the extended set of benchmark (K) and test (N) assets.

In our analysis, the K benchmark assets are represented by six of the seven main tax categories, while the N test assets are represented by the seventh remaining tax (i.e. $K = 6$ and $N = 1$). MVS analysis allows us to examine if the addition of the seventh tax to our portfolio of six taxes shifts the mean-variance efficient frontier. In other words, we test if including the seventh tax category in the portfolio of taxes improves the risk-return trade-off relative to the original benchmark of six taxes. We replicate this analysis with each of the seven tax categories acting as the test asset (N) in turn.

In practice, the typical spanning test involves the regression of the N test asset returns on the returns of the K benchmark assets as follows:

$$R_{N,t} = \alpha + \beta R_{K,t} + \varepsilon_t \qquad (4)$$

with $\varepsilon_t \sim N(0, \Sigma)$, $\alpha = E[R_{N,t}] - \beta E[R_{K,t}] = \mu_N - \beta\mu_K$ and $\beta = V_{NK}V_{KK}^{-1}$. In the same way that the returns of financial assets represent the price growth of those assets, our analysis uses the "returns" for each of the tax categories, or the revenue growth rate for each tax.

In defining $\delta = 1_N - \beta 1_K$, Huberman and Kandel (1987) and Kan and Zhou (2008) provide the necessary and sufficient conditions for spanning in terms of a restriction on α and δ such that we can test the following null hypotheses:

$$H^1_0\!: \alpha = 0_N, \quad H^2_0\!: \delta = 0_N \qquad (5)$$

It follows that a test of whether the benchmark assets (K) span the extended set of benchmark plus test assets (K+N) involves jointly testing the above hypotheses. A failure to reject the null hypothesis implies that for the test asset (N), it is possible to form a portfolio of the benchmark assets (K), that has the same expected return (because $\alpha = 0_N$ and $\beta 1_K = 1_N$) but a lower variance (as $R_{K,t}$ and ε_t are uncorrelated while $Var\ (\varepsilon_t)$ is positive definite).

From an empirical standpoint, we use Ordinary Least Squares (OLS) estimation, and carry out a Wald test of the restrictions in (5). A detailed derivation of the statistical test of spanning (adapted from Kan and Zhou, 2008; Berrill and Kearney, 2010) is contained in the Appendix of this paper.

VI EMPIRICAL ANALYSIS AND RESULTS

6.1 Correlation and Cointegration of Tax Revenue

Modern Portfolio Theory advances that diversification benefits arise from holding a portfolio of underlying assets that do not display strong and positive linkages.

Table 2 displays bivariate correlation coefficients for all pairs of the seven tax categories. This shows the strength of the linear relationship between each pair of revenue streams. Overall, Customs is the least correlated with other taxes generally, while VAT appears to be the most correlated. The highest individual correlation is between VAT and Excise Duty (0.78), as expected, given that these taxes are levied on some of the same products. Similarly, Stamp Duty and Capital Taxes have the second highest correlation (0.67), reflective of the fact that they (like VAT and Excise Duty) share elements of the same tax base.

Table 2: Correlation Matrix of Tax Categories, 1985-2017

	Income	*VAT*	*Excise*	*Corporation*	*Stamp*	*Customs*	*Capital*
Income	*1.00*	0.47	0.34	0.37	0.46	0.27	0.46
VAT	0.47	*1.00*	0.78	0.47	0.59	0.31	0.60
Excise	0.34	0.78	*1.00*	0.44	0.54	0.19	0.50
Corporation	0.37	0.47	0.44	*1.00*	0.12	0.01	0.44
Stamp	0.46	0.59	0.54	0.12	*1.00*	0.22	0.67
Customs	0.27	0.31	0.19	0.01	0.22	*1.00*	0.15
Capital	0.46	0.60	0.50	0.44	0.67	0.15	*1.00*

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank.
Notes: The correlation coefficient is bound above by 1 (indicating perfect positive correlation), and below by -1 (indicating perfect negative correlation).

We follow this correlation analysis by examining the long-term linkages among the seven tax categories. We model these linkages as cointegrating relations using the Johansen-Juselius test (1990). As previously discussed, this is a commonly used method in assessing the level of integration within a system or between a pair of variables, and has been employed extensively in studies of financial market integration (for example, see Quayes and Jamal, 2016; Neaime, 2016; and Babalos *et al.,* 2016).

Figure 4 shows the plot of the scaled trace statistic for each sub-sample in a recursive framework (scaled at the 90 per cent significance level). When the graph trends above the horizontal line, there is cointegration among the tax revenue streams. As shown, there is considerable intertemporal instability in the level of cointegration among the tax revenues. We find particularly strong evidence of cointegration as the sample expands to include tax revenue from 2004 to 2008, suggesting a more integrated tax system in the years prior to the crisis. By the end of the sample, we observe relatively weaker evidence for cointegration, with a decline in the level of the scaled trace statistic between 2008 and 2011 (in line with fiscal consolidation efforts in the aftermath of the crisis) and between 2015 and 2017 (potentially reflecting the level shift in Corporation Tax due to multinationals relocating their assets to Ireland).

Figure 4: Plot of the Scaled Trace Statistics

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank.
Notes: The dashed line shows the scaled trace statistic (at the 90 per cent significance level). When the dashed line is above the solid horizontal line, there is cointegration in the system. The 90 per cent critical value is 118.5. The graph begins at 2000, as the initial estimation window runs from 1984 to 2000, covering the first 16 years of the sample. The sample is then expanded incrementally by a single year, and the model is re-estimated over 1984-2001. This process is repeated until the end of the sample is reached (in 2017), at which point the analysis is over the entire sample (from 1984-2017).

Generally, towards the end of the sample, the level of cointegration remains substantially below pre-crisis levels. This suggests that, while substantial linkages between the tax revenue streams remain, these linkages are weaker than in the pre-crisis period, and could be explained by efforts to broaden the tax base in response to the fiscal crisis. This merits further analysis and consideration in future work, perhaps involving the use of an adjusted revenue series that disentangles the impact of policy change.

6.2 Estimating Optimal Revenue Shares

In this section we estimate optimal revenue shares as those that comprise a portfolio that lies on the mean-variance efficient frontier. Generally, the efficient frontier displays the set of portfolios that offer the highest expected return for a given level of risk, or the lowest risk for a given level of expected return. Portfolios that lie beyond the frontier are unattainable, and are not part of the investor's investment opportunity set, while those that lie below the frontier are sub-optimal, as it is possible to select an alternative portfolio that offers a higher expected return for the same level of risk.

In constructing the efficient frontier, we begin by establishing the global minimum variance portfolio, that is the combination of shares that achieves the lowest possible level of risk (the minimum variance portfolio marks the beginning of the frontier). This involves optimising (minimising) the portfolio variance equation in (1) subject to the constraint that the share of taxes in the portfolio be non-negative and sum to 1.

In doing so, we establish the variance minimising shares detailed in Column 2 of Table 3. As shown, the variance minimising portfolio of taxes would include a majority share of Excise Duty (65 per cent) followed by Income Tax (33 per cent) with some additional revenue from Customs (2 per cent). Although the inclusion of Customs appears unusual from a risk-return perspective, this potentially reflects the relatively low level of correlation between Customs and each other tax. These shares differ substantially from the historical average for the sample, and from the 2017 shares, also detailed in Table 3.

The global minimum variance portfolio shows how the tax system would be structured if policymakers were motivated solely by the goal of minimising Exchequer revenue volatility. However, targeting this global minimum variance portfolio may not be desirable or feasible from a policymaker's perspective. Specifically, the underlying base may not generate sufficient revenue, and there are important issues of progressivity and equity that may make a particular allocation undesirable. Furthermore, governments may not have the political capital or will to make substantial changes to the current tax system. We consider this first estimation as a theoretical benchmark exercise.

For a more realistic outcome, we replicate this analysis imposing minimum and maximum "feasibility constraints" that are based on more reasonable assumptions.

Table 3: Variance Minimising Versus Actual Shares for Each Tax Category

Tax Category	*Variance Minimising*	*Actual average (1985-2017)*	*2017*	*Sample Min.*	*Sample Max.*	*Feasibility Constraints*
Income Tax	33%	37%	40%	27%	44%	42%
VAT	0%	28%	27%	25%	33%	25%
Excise Duty	65%	17%	12%	12%	25%	25%
Corporation Tax	0%	11%	16%	4%	16%	4%
Stamp Duty	0%	4%	2%	2%	8%	2%
Customs	2%	1%	1%	0%	2%	2%
Capital Taxes	0%	2%	3%	1%	8%	1%

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank.
Notes: The above figures are rounded to the nearest percentage, which may affect totals.

We impose minimum and maximum limits on what each share can be, that are intended to more accurately represent the limitations on the opportunity set of policymakers. These limits are defined by the historical sample minimum and maximum shares for each tax, and are as shown in Columns 5 and 6 of Table 3. The variance minimising shares subject to feasibility constraints are shown in the final column. As can be seen, Income Tax, Excise and Customs converge to their sample maximum values, while VAT, Corporation Tax and Capital Taxes converge to their sample minimum. Comparing the minimum variance portfolio under "feasibility" constraints to the tax mix as of 2017, Corporation Tax appears to be over-weighted as of 2017 (by 12 per cent), while Excise Duty appears to be under-weighted (by 13 per cent). There are minimal differences for the remainder of the taxes. The difference between the two portfolios, in terms of risk and return, is further explored in Figure 5.

These volatility minimising shares, with and without feasibility constraints, represent useful benchmark portfolios. These shares could inform incremental policy change towards reducing Exchequer revenue volatility, or could inform alternative strategies to determining the amount of revenue to be contributed to a counter-cyclical fiscal fund (or Rainy Day Fund). The idea of linking Rainy Day Fund allocations to short-term volatility in tax revenue is explored in more detail later, when discussing the practical implications of the paper for Irish fiscal policymaking.

Garrett (2006) states that, as the variance minimising shares are a function of the sample period used, the stability of these estimates should be examined over multiple periods of the business cycle. In effect, this involves establishing the minimum variance portfolio for each sub-period (absent any feasibility constraints), and noting the optimal tax shares in each case. We examine the sub-period from

1985 to 1999, and also examine pre-crisis (2000-2007) and post-crisis (2008-2017) sub-periods. These results are shown in Table 4.

We find that the variance minimising share of Income Tax has been declining over time, from 32 per cent between 1985 to 1999, to 20 per cent in the pre-crisis period, and 16 per cent in the post-crisis period. Conversely, the variance minimising share of Excise Duty has increased from 41 per cent pre-crisis, to 84 per cent post-crisis. While VAT and Corporation Tax have variance minimising shares of 30 per cent and 10 per cent pre-crisis respectively, these shares drop to 0 per cent post-crisis. In optimally allocating shares to minimise portfolio variance, it appears that there has been a move away from a more diversified tax portfolio, towards greater concentration on less volatile taxes (as also observed in Figure 3).

Table 4: Variance Minimising Shares for Each Tax Category, Pre- and Post-Crisis

Tax Category	*Full Sample 1985-2017*	*1985-1999*	*Pre-Crisis 2000-2007*	*Post-Crisis 2008-2017*
Income Tax	33%	32%	20%	16%
VAT	0%	0%	30%	0%
Excise Duty	65%	56%	41%	84%
Corporation Tax	0%	3%	10%	0%
Stamp Duty	0%	0%	0%	0%
Customs	2%	9%	0%	0%
Capital Taxes	0%	0%	0%	0%

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank.
Notes: The above figures are rounded to the nearest percentage which may affect totals.

Having established the global minimum variance portfolio, we proceed by constructing the efficient frontier. We optimise Equation (1) for different levels of risk/volatility, and plot the results in risk-return space as shown in Figure 5. We also plot each of the portfolio allocations detailed in Table 3, that is:

- the global minimum variance portfolio;
- the minimum variance portfolio subject to feasibility constraints (the feasible allocation);
- the portfolio consisting of sample average shares; and,
- the portfolio consisting of 2017 shares.

As shown, with the exception of the global minimum variance portfolio (which by definition lies on the efficient frontier), all of the plotted portfolios lie below the frontier. This indicates that, in risk-return terms, these are sub-optimal portfolios, and the tax shares could be reallocated in a way that improves the level of expected return for the inherent level of risk in each case.

Figure 5: Efficient Frontier of the Complete Set of Taxes, 1985-2017

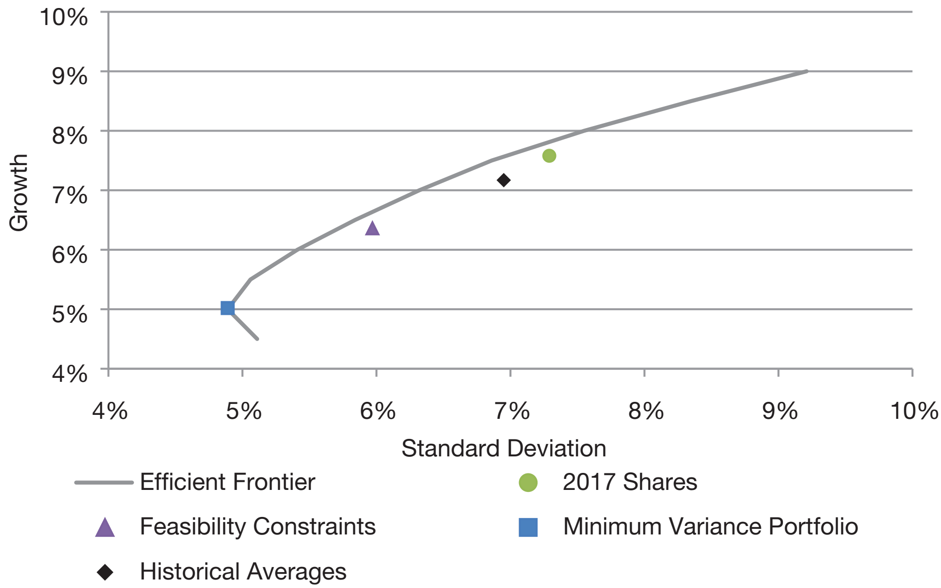

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank.

6.3 Mean-Variance Spanning

In the preceding section, we established the efficient frontier that illustrates the optimal portfolio allocations for the complete set of seven taxes. We proceed with a mean-variance spanning analysis; that is for each possible combination of six of the seven tax categories, we assess if the addition of the seventh tax shifts the mean-variance efficient frontier. In effect, this informs us whether the addition of the seventh tax in each case presents diversification benefits (i.e. improvements in the risk-return trade-off).

Figures 6a to 6f contain the efficient frontiers for each combination of six taxes. As shown, there is a visual improvement in the efficient frontier from adding Income Taxes to a portfolio consisting of the six other tax categories, with a visible frontier shift. We observe only minor improvements from including Corporation Tax, VAT and Excise Duty, while the frontiers with and without Stamp Duty and Customs, appear to overlap (or "span").

While this graphical analysis provides a useful and intuitive interpretation of the benefits from adding the additional tax in each case, we proceed with a test of the statistical significance of these shifts in the efficient frontiers.

Figure 6a: Income Tax

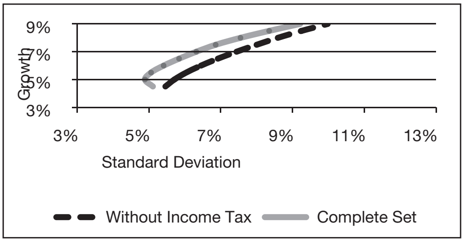

Figure 6b: VAT

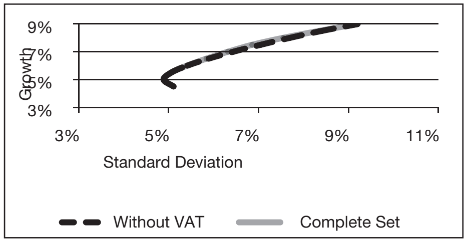

Figure 6c: Corporation Tax

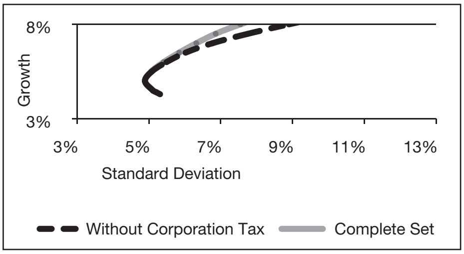

Figure 6d: Excise Duty

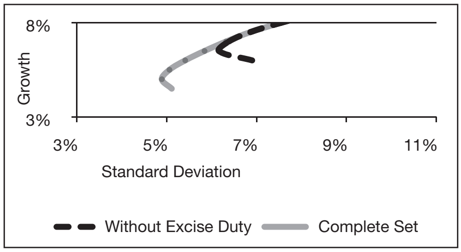

Figure 6e: Stamp Duty

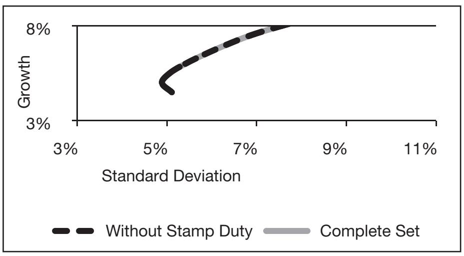

Figure 6f: Customs

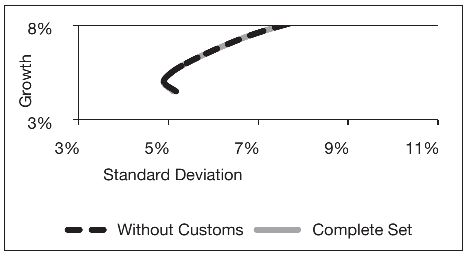

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank.

6.4 Statistical Tests of Spanning

The results of our statistical tests of mean-variance spanning are shown in Table 5. Intuitively, the lower the p value (in parentheses), the higher the expected return per unit of risk from adding the additional tax category to the State's portfolio of taxes (or the lower the risk per unit of return). As shown, we can reject the null hypothesis of spanning at the one per cent level for both Income Tax and Excise Duty. Consistent with our graphical analysis of the efficient frontiers in Figures 6a and 6d, this suggests that there are benefits in terms of the risk-return trade-off from including Income Tax and Excise Duty, to a portfolio consisting of the six other taxes.

However, this result does not hold for the rest of the tax categories, that is, there are no statistically significant benefits in terms of the risk-return trade-off, from adding VAT, Corporation Tax, Stamp Duty, Customs or Capital Taxes to a portfolio of the six other taxes. This is a somewhat surprising result, particularly as VAT offers a better risk-return trade-off (i.e. a higher ratio of growth to standard deviation) than Excise Duty. However, this could be explained by the relatively high correlation between the two (at 0.78, see Table 2) and the lower correlation between Excise Duty and each other tax category generally, relative to VAT (a lower correlation allows for potentially greater diversification benefits). In addition, VAT offers a similar average growth rate to Income Tax (6.7 per cent versus 6.6 per cent respectively) but with a higher level of volatility (7.9 per cent to 6.9 per cent respectively).

Table 5: Mean Variance Spanning Results

Test Assets	*Test Statistic*	*(p value)*
Income Tax	6.39***	*(0.006)*
VAT	0.70	*(0.504)*
Excise Duty	14.40***	*(0.000)*
Corporation Tax	1.86	*(0.175)*
Stamp Duty	1.35	*(0.278)*
Customs	0.10	*(0.908)*
Capital Taxes	0.32	*(0.726)*

Source: Authors' analysis based on Department of Public Expenditure & Reform – Databank.
Notes: In each case, the null hypothesis (H_0) is that the benchmark assets span the extended set of benchmark plus test assets. The test assets are as indicated in each case, while the benchmark assets are the remaining six tax categories. *** Indicates statistical significance at the 1 per cent level. Capital Taxes refers to both Capital Acquisitions Tax (CAT) and Capital Gains Tax (CGT). Outturns for CGT and CAT were reported jointly under Capital Taxes until 2000).

We can conclude that there is only a statistically significant shift in the mean-variance efficient frontier from adding holdings of Income Tax or Excise Duty, to a State portfolio consisting of the six remaining tax categories. Whereas, for VAT, Corporation Tax, Stamp Duty, Customs and Capital Taxes, it is possible to form a tax portfolio of the six remaining taxes that has the same expected return but a lower volatility.

VII CONCLUSION

7.1 Summary of Results

This paper presents several key empirical findings. We establish that a trade-off exists between tax revenue growth and volatility, more typically associated with the returns of financial assets: the most volatile taxes also experience the highest growth rate on average.

From a diversification perspective, we find that the overall volatility of the State's tax portfolio moves in line with the level of diversification across the seven main taxes (as indicated by the HHI measure). In the period preceding the crisis, while the tax portfolio was relatively more diversified, it was also more volatile, with greater holdings of more volatile revenue streams such as Stamp Duty and Capital Taxes. Alongside the fiscal consolidation that followed the 2008 economic and fiscal crisis, we find that the volatility of the tax portfolio declined, as the portfolio became more concentrated around less volatile taxes (e.g. Income Tax). We further establish strong (but time-varying) cointegration among the tax revenue streams, implying that there are minimal diversification benefits generally.

Following this, we establish the portfolio that minimises overall portfolio volatility (that is, the minimum variance portfolio), and find that it consists of a majority holding of Excise Duty (65 per cent), with some Income Tax (33 per cent) and a small allocation to Customs (2 per cent). To better represent the options available to policymakers, we calculate the minimum variance portfolio subject to feasibility constraints. These constraints are equal to the sample minimum and maximum bounds on the size of each tax share in the tax portfolio. Comparing these more realistic minimum variance allocations to the tax portfolio as of 2017, Corporation Tax appears to be over-weighted as of 2017 (by 12 per cent), while Excise Duty appears to be under-weighted (by 13 per cent), with minimal differences for the remainder of the taxes.

In constructing the efficient frontier for the State's tax portfolio, we find that the 2017 allocation is sub-optimal from a mean-variance efficiency perspective. That is, an alternative combination of the seven taxes would have allowed for a greater portfolio return for the same level of volatility. Finally, in terms of mean-variance spanning, we observe a statistically significant shift in the efficient frontier from adding Income Tax and Excise Duty to a benchmark portfolio consisting of the six remaining taxes, while this result does not hold for each of the other taxes.

Our paper estimates volatility minimising tax revenue shares from the perspective of Modern Portfolio Theory. However, the resulting shares may be undesirable or infeasible from the policymaker's perspective, when also considering issues of equity and efficiency. This paper aims to make a novel contribution by providing a theoretical discussion of how the State's portfolio of taxes might be structured if guided solely by the principles of risk and return. Furthermore, practical policy implications of contemporary relevance to Irish fiscal policy follow from this analysis.

7.2 Policy Implications

Exchequer revenue volatility can complicate prudential fiscal planning, and risks undermining the stability and sustainability of the public finances. Generally, volatile revenues are harder to predict, and come with sizable forecast errors (see Hannon *et al.*, 2015). The volatility of Corporation Tax receipts in particular, has been the focus of contemporary debate on fiscal policy in Ireland. For ease of fiscal planning, and to ensure that public spending of a permanent nature is linked to stable sources of revenue, the State might be interested in pursuing incremental policy change that re-balances the tax system to reduce Exchequer revenue volatility. The State may also wish to allocate revenue surprises (windfall receipts) to a counter-cyclical fiscal fund (or Rainy Day Fund). This would assist in avoiding using transient revenues to fund permanent expenditure. Our findings have some important policy implications that relate to these issues.

First, in designing tax policies that alter the composition of State tax revenue, consideration should be given to the implications for volatility and growth of the overall tax portfolio (and the risk-return trade-off with respect to the taxes in that portfolio). While certain revenue streams such as Capital Taxes and Stamp Duties have a high growth potential, particularly when the economy is growing fast, these are also highly volatile, and in a recession they may decline rapidly. We find that the volatility of the Irish tax portfolio was lower when the portfolio was concentrated around relatively less volatile revenues (e.g. Income Tax). Furthermore, benefits (in terms of lower portfolio volatility) from diversification across the seven tax categories are limited due to strong cointegration.

Second, in the event that the tax portfolio is over-weighted with relatively more volatile taxes, we emphasise the importance of a counter-cyclical fiscal fund in mitigating the risks of an over-reliance on potentially unsustainable revenues. With the stated goal of implementing counter-cyclical fiscal policy, the Irish Government is establishing a Rainy Day Fund (RDF). At present, pre-committed annual Exchequer contributions of €500 million have been announced for allocation to the RDF for 2019 through 2023 (in addition to €1.5 billion in seed capital from the Irish Strategic Investment Fund). The RDF's establishing legislation also outlines that one-off or windfall revenues (mostly related to unanticipated levels of Corporation Tax receipts) may be transferred to the RDF.

RDFs can be used to stimulate the economy in a downturn, effectively smoothing the business cycle. Draw-downs from an RDF can provide a continued source of funds for key public services, particularly at a time of increasing demand. We have seen evidence of similar funds being used in the Irish context, with the National Training Fund (sourced by a levy on employers and used to fund certain further education and training programmes) and the Social Insurance Fund effectively acting as counter-cyclical funds. While payments into these Funds are inherently pro-cyclical (larger in times of high employment), surpluses generated in periods of strong economic performance allowed for a relatively stable source of funding throughout the fiscal crisis.

Our analysis presents a potential tool to guide allocations to such a fund. This would involve allocating revenues based on the extent to which the share of a tax exceeds, in a given year, an estimated benchmark share. This benchmark share could reflect the volatility minimising share. For a more reasonable scenario, in light of the constraints on policymakers, the benchmark could be based on the volatility minimising share subject to certain feasibility constraints. These constraints may be determined by the minimum and maximum share bounds for the tax, over a defined or rolling time period.[2] As an example, consider the temporarily large share of Stamp Duty during the Irish housing boom of the early-mid 2000s, or the more recent surge in Corporation Tax revenues, as shown in Figure 2. This approach would help the State to avoid spending transient or windfall tax revenues, and would assist in linking public expenditure to more sustainable sources of revenue.

Finally, policymakers may use these variance minimising shares as a benchmark to inform incremental (rather than radical) changes to tax policy. Generally, the tax shares estimated in this paper should be seen as indicative (rather than prescriptive), informed by the historical risk-return profile for each tax.

7.3 Future Work

Future work might replicate elements of our analysis on an adjusted revenue series, removing the discretionary component from tax revenue volatility (i.e. controlling for the impact of policy change). This might involve using the proportional adjustment method to disentangle the impact of discretionary measures (see Princen *et al.*, 2013 and Barrios and Fargnoli, 2010, for an analysis of discretionary tax measures in the European Union; and Casey and Hannon, 2016, for an analysis of Irish Corporation Tax specifically). To date, this method has only been applied in an analysis of tax elasticities, and the literature applying Modern Portfolio Theory to tax revenue volatility has not made this adjustment. As such, this would represent a novel contribution to the literature.

Additional work in this area could examine statistical tests of mean-variance spanning in a sub-period framework, across different time periods and at different points of the business cycle. Future work might also seek to explore the origins of the risk-return trade-off identified in our paper, and the drivers of revenue volatility for each of the seven taxes more broadly. This in turn would make a valuable contribution in terms of the accuracy of revenue forecasts, and would aid in fiscal planning.

[2] For example, the benchmark shares could be determined by estimating a rolling five-year minimum variance portfolio (subject to feasibility constraints), to account for structural changes in the tax system or in the wider macro-economy over time. When a share exceeds this benchmark, allocation to the RDF would be considered. Therefore, at any point, the policymaker would examine what the volatility minimising portfolio would consist of, based on the risk-return behaviour of the taxes over the preceding five years only (subject to some lower and upper bounds on the shares).

REFERENCES

Acheson, J., D. D. Yota, D. Lambert and E. L. W. Morgenroth, 2017. "Income Tax Revenue Elasticities in Ireland: An Analytic Approach", *Research Series*, No. 59, Economic and Social Research Institute.

Addison-Smyth, D. and K. McQuinn, 2010. "Quantifying Revenue Windfalls from the Irish Housing Market", *The Economic and Social Review*, Vol. 41, No. 2, pp. 201-223.

Addison-Smyth, D. and K. McQuinn, 2016. "Assessing the Sustainable Nature of Housing-Related Taxation Receipts: The Case of Ireland", *Journal of European Real Estate Research,* June.

Babalos, V., G. M. Caporale and N. Spagnolo, 2016. "Equity Fund Flows and Stock Market Returns in the US Before and After the Global Financial Crisis: A VAR-GARCH-In-Mean Analysis", *DIW Berlin Discussion Paper*, No. 1583.

Barrios, S. and R. Fargnoli, 2010. "Discretionary Measures and Tax Revenues in the Run-up to the Financial Crisis", *European Economy Economic Papers*, 419.

Berrill, J. and C. Kearney, 2010. "Firm-Level Internationalisation and the Home Bias Puzzle", *Journal of Economics and Business*, Vol. 62, No. 4.

Casey, E. and A. Hannon, 2016. "Challenges Forecasting Irish Corporation Tax", *Analytical Note No. 10*, Irish Fiscal Advisory Council.

Cornia, G. C. and R. D. Nelson, 2010. "State Tax Revenue Growth and Volatility", *Regional Economic Development*, Vol. 6, No. 1, pp. 23-58, Federal Reserve Bank of St. Louis.

Crain, M. W., 2003. *Volatile States: Institutions, Policy, and the Performance of American State Economies*, University of Michigan Press.

Deli, D. Y., A. G. Rodriguez, I. Kostarakos and P. Varthalitis, 2018. "Dynamic Tax Revenue Buoyancy Estimates for a Panel of OECD Countries", *ESRI Working Paper*, No. 592, Economic and Social Research Institute.

Eiling, E., B. Gerard, P. Hillion and F. A. de Roon, 2012. "International Portfolio Diversification: Currency, Industry and Country Effects Revisited", *Journal of International Money and Finance*, Vol. 31, No. 5, pp. 1249-1278.

Garrett, T. A., 2006. "Evaluating State Tax Revenue Variability: A Portfolio Approach", *Working Paper Series*, 2006-008A, Federal Reserve Bank of St. Louis.

Hannon, A., 2014. "Tax Forecasting Error Decomposition", *Analytical Note*, No. 3, Irish Fiscal Advisory Council.

Hannon, A., E. Leahy and R. O'Sullivan, 2015. "An Analysis of Tax Forecasting Errors in Ireland", *Working Paper*, No. 3, Irish Fiscal Advisory Council.

Hansen, H. and S. Johansen, 1993. "Recursive Estimation in Cointegrated VAR-Models", *Discussion Papers*, No. 92-13, University of Copenhagen Department of Economics.

Huberman, G. and S. Kandel, 1987. "Mean-variance Spanning", *The Journal of Finance*, Vol. 42, No. 4.

Jobson, J. D. and B. Korkie, 1989. "A Performance Interpretation of Multivariate Tests of Asset Set Intersection, Spanning, and Mean-Variance Efficiency", *Journal of Financial and Quantitative Analysis*, Vol. 24, No. 2, pp. 185-204.

Johansen, S. and K. Juselius, 1990. "Maximum Likelihood Estimation and Inference on Cointegration – With Applications to the Demand for Money", *Oxford Bulletin of Economics and Statistics*, Vol. 52, No. 2.

Johansson, Å., 2016. "Public Finance, Economic Growth and Inequality: A Survey of the Evidence", *OECD Economics Department Working Papers*, No. 1346, OECD Publishing, Paris, https://doi.org/10.1787/094bdaa5-en.

Kan, R. and G. Zhou, 2008. "Tests of Mean-Variance Spanning", *OLIN Working Paper*, No. 99-05.

Lucey, B. M. and S. Voronkova, 2005. "Russian Equity Market Linkages Before and After the 1998 Crisis: Evidence from Time-Varying and Stochastic Cointegration Tests", *BOFIT Discussion Paper*, 12/2005.

Markowitz, H., 1952. "Portfolio Selection", *The Journal of Finance*, Vol. 7, No. 1.

Matthews, J., 2005. "Tax Revenue Volatility and a State-wide Education Sales Tax", *Fiscal Research Centre*, Policy Brief Number 109, Andrew Young School of Policy Studies.

McQuinn, K. and J. M. Roche, 2016. "Efficient Frontiers and Fiscal Stability: An Ex-ante and Ex-post Application to the Irish Public Finances", *ESRI Working Paper*, No. 538, Economic and Social Research Institute.

Neaime, S., 2016. "Financial Crises and Contagion Vulnerability of MENA Stock Markets", *Emerging Markets Review*, Vol. 27, pp. 14-35.

O'Hagan-Luff, M. and J. Berrill, 2015. "Why Stay-at-Home Investing Makes Sense", *International Review of Financial Analysis*, Vol. 38, pp. 1-14.

Princen S., G. Mourre, P. Paternoster and G-M. Isbasoiu, 2013. "Discretionary Tax Measures: Pattern and Impact on Tax Elasticities", *European Economy Economic Papers*, 499.

Quayes, S. and A. M. M. Jamal, 2016. "Impact of Demographic Change on Stock Prices", *The Quarterly Review of Economics and Finance*, Vol. 60, pp. 172-179.

Revenue Commissioners, 2017. "Income Tax and Corporation Tax Distribution Statistics".

Revenue Commissioners, 2018. "Corporation Tax 2017 Payments and 2016 Returns".

APPENDIX: DERIVATION OF THE SPANNING TESTS

Tests of spanning can be derived as follows (adapted from Kan and Zhou, 2008; Berrill and Kearney, 2010). Consider the following model in matrix notation (note that this is simply Equation (4) in matrix form):

$$R = X\beta + \Sigma, \tag{6}$$

The unconstrained maximum likelihood estimates of β and Σ are determined by

$$\hat{\beta} = (X'X)^{-1}(X'R) \text{ and, } \hat{\Sigma} = \frac{1}{T}(R - X\hat{\beta})'(R - X\hat{\beta}), \tag{7}$$

The spanning test can be derived by defining:

$$\hat{\mu} = \sum_{t=1}^{T} \frac{R_t}{T} \text{ and } \hat{V} = \frac{1}{T}\sum_{t=1}^{T} (R_t - \hat{\mu})(R_t - \hat{\mu})',$$

and four constants *a, b, c* and *d*. Consider the two efficient frontiers: the frontier for *K* assets, and the frontier for *K*+*N* assets. Firstly, for the *K* assets:

$$\hat{a}_K = \hat{\mu}'_K \hat{V}^{-1}{}_{11} \hat{\mu}_K \quad \hat{b}_K = \hat{\mu}'_K \hat{V}^{-1}{}_{11} 1_K \quad \hat{c}_K = \hat{1}'_K \hat{V}^{-1}{}_{11} 1_K \quad \hat{d}_K = \hat{a}_K \hat{c}_K - \hat{b}^2{}_K$$

While for the *K*+*N* assets:

$$\hat{a}_{K+N} = \hat{\mu}'_{K+N} \hat{V}^{-1} \hat{\mu}_{K+N} \qquad \hat{b}_{K+N} = \hat{\mu}'_{K+N} \hat{V}^{-1} 1_{K+N}$$

$$\hat{c}_{K+N} = \hat{1}'_{K+N} \hat{V}^{-1} 1_{K+N} \qquad \hat{d}_{K+N} = \hat{a}_{K+N} \hat{c}_{K+N} - \hat{b}^2{}_{K+N}$$

In moving to the frontier of the extended set of assets, *K*+*N,* from the frontier of the benchmark assets, *K,* the above constants change by:

$$\Delta\hat{a} = \hat{a}_{K+N} - \hat{a}_K \qquad \Delta\hat{b} = \hat{b}_{K+N} - \hat{b}_K \qquad \Delta\hat{c} = \hat{c}_{K+N} - \hat{c}_K$$

We can then form the following matrices:

$$\hat{G} = \begin{vmatrix} 1 + \hat{a}_K & \hat{b}_K \\ \hat{b}_K & \hat{c}_K \end{vmatrix} \text{ and, } \hat{H} = \begin{vmatrix} \Delta\hat{a} & \Delta\hat{b} \\ \Delta\hat{b} & \Delta\hat{c} \end{vmatrix} \tag{8}$$

Combining these two matrices, allowing $\hat{\Sigma}$ to denote the unconstrained maximum likelihood estimate of Σ with *K*+*N* assets in (7), and allowing $\tilde{\Sigma}$ to denote the constrained maximum likelihood estimate of Σ with *K* assets in (7), and defining

$U = |\hat{\Sigma}\ \tilde{\Sigma}^{-1}|$, then the likelihood ratio test of whether the K benchmark assets span the extended set of $K+N$ benchmark and test assets is given by:

$$LR = -\,T\ ln\,(U) \tag{9}$$

where,

$$U = |\hat{\Sigma}\ \tilde{\Sigma}^{-1}| = \frac{|\hat{G}|}{|\hat{G} + \hat{H}|} = \frac{(1+\hat{a}_K)\hat{c}_K - \hat{b}^2{}_K}{(1+\hat{a}_{K+N})\hat{c}_{K+N} - \hat{b}^2{}_{K+N}} = \left(\frac{\hat{c}_K}{\hat{c}_{K+N}}\right)\left(\frac{1+\dfrac{\hat{d}_K}{\hat{c}_K}}{1+\dfrac{\hat{d}_{K+N}}{\hat{c}_{K+N}}}\right) \tag{10}$$

Under the null hypothesis, the distribution of the likelihood ratio test is:

$$F = \left(\frac{T-K-N}{N}\right)\left(U^{-½}-1\right) = \left(\frac{T-K-N}{N}\right)\left[\left(\frac{\sqrt{\hat{c}_{K+N}}}{\sqrt{\hat{c}_K}}\right)\left(\frac{\sqrt{1+\dfrac{\hat{d}_{K+N}}{\hat{c}_{K+N}}}}{\sqrt{1+\dfrac{\hat{d}_K}{\hat{c}_K}}}\right) - 1\right] \tag{11}$$

as demonstrated by Huberman and Kandel (1987) and Jobson and Korkie (1989). The ratio given by $\left(\frac{\sqrt{\hat{c}_K}}{\sqrt{\hat{c}_{K+N}}}\right)$ is the ratio of the standard deviations of the minimum variance portfolios of the K benchmark assets and the $K+N$ extended set of benchmark and test assets, bound below by 1. The ratio given by $\left(\frac{\sqrt{1+\frac{\hat{d}_{K+N}}{\hat{c}_{K+N}}}}{\sqrt{1+\frac{\hat{d}_K}{\hat{c}_K}}}\right)$ is the length of the asymptote to the $K+N$ efficient frontier divided by the equivalent to the restricted frontier of the K benchmark assets, bound below by 1, as shown by Kan and Zhou (2008).

While Huberman and Kandel (1987) suggest testing the imposed constraints in (5) using the likelihood ratio (LR) test, Kan and Zhou (2008) demonstrate that the likelihood ratio (LR), Lagrange multiplier (LM), and Wald (W) tests are closely related spanning tests (see visualisation in Figure A.1). Indeed, Kan and Zhou (2008) show that for the case when $N = 1$, the Wald test is the most powerful of the three tests.

Following this, we use Ordinary Least Squares estimation, and carry out a Wald test of the restrictions in (5).

Figure A.1: Geometry of the MVS Tests

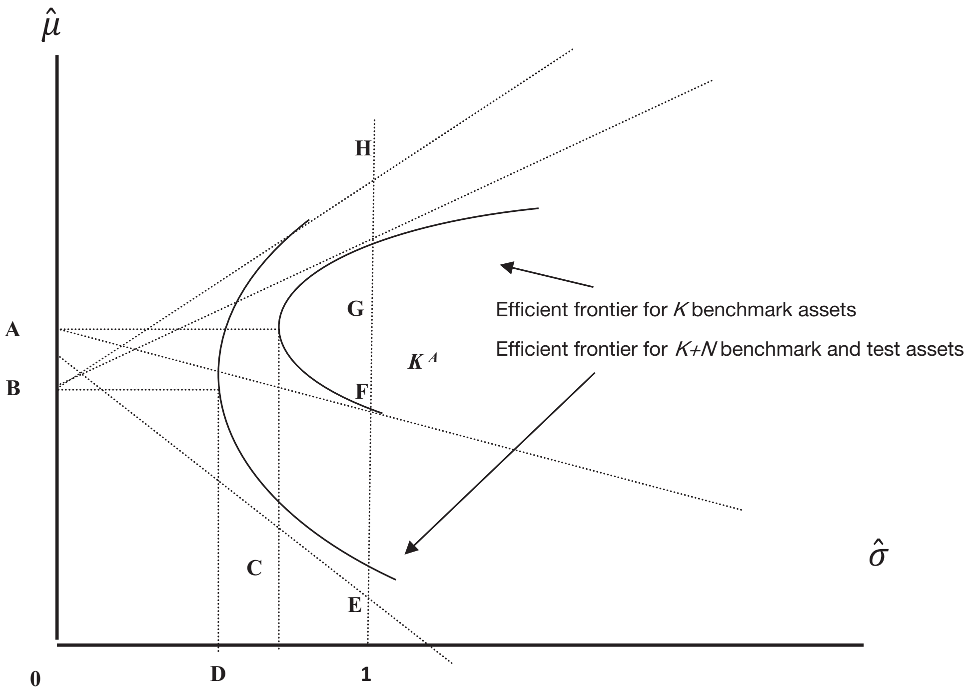

Source: Based on Kan and Zhou (2008) and Berrill and Kearney (2010).
Notes: The geometry of the Wald (W), Lagrange Multiplier (LM), and Likelihood Ratio (LR) spanning tests, taken from Kan and Zhou (2008) and Berrill and Kearney (2010), are as below:

$$W = \left(\frac{OC}{OD}\right)^2 - 1 + \left(\frac{AE}{AF}\right)^2 - 1$$

$$LM = 1 - \left(\frac{OD}{OC}\right)^2 + 1 - \left(\frac{BG}{BH}\right)^2$$

$$LR = \left(\frac{OC}{OD}\right)\left(\frac{BH}{AF}\right) - 1$$